TestOut®

TestOut Security Pro - English 8.0.x

ACCESSIBILITY GUIDE

Modified 2023-11-15

Powered by LABSIM

Security Pro – v8.0.x
Accessibility Document

NOTE: This document is to be used only by authorized instructors and students who have accessibility needs. Other use or distribution is prohibited.

1.0 Security Concepts

1.1 Security Introduction

As you study this section, answer the following questions:

- What is information security?
- What challenges does a security professional face?
- What aspects create a proactive approach to security?
- What is the difference between integrity and non-repudiation?
- What are the three main goals of the CIA of Security?
- What are the key components of risk management?
- What are the three types of threat agents?

The key terms for this section include:

Term	Definition
Security operations center (SOC)	The location where security professionals monitor and protect critical information assets in an organization.
Development and operations (DevOps)	A combination of software development and systems operations and refers to the practice of integrating one discipline with the other.
DevSecOps	A combination of software development, security operations, and systems operations and refers to the practice of integrating each discipline with the others.
Computer incident response team (CIRT)/computer security incident response team (CSIRT)/computer emergency response team (CERT)	Team with responsibility for incident response. The CSIRT must have expertise across a number of business domains (IT, HR, legal, and marketing, for instance).

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.1 Compare and contrast various types of security controls.
	• Categories
	  ○ Technical
	  ○ Managerial
	  ○ Operational

	○ Physical 1.2 Summarize fundamental security concepts. • Confidentiality, Integrity, and Availability (CIA) 2.1 Compare and contrast common threat actors and motivations. • Threat actors ○ Nation-state ○ Organized crime

1.1.1 The Security Landscape (Lesson Video)

Transcript:

Welcome to TestOut's Security Pro course. This course is designed to help you understand the Information security landscape and will prepare you to become a security professional.

But before we can dive into the ever-changing murky waters that is information security, we first need to talk about the security landscape that all security professionals must face.
In today's world, cyber criminals are a very real and dangerous threat. One way to think about Information security is that of a never-ending arms race, with the sophistication of the weapons being used advancing exponentially. Every day cyber criminals are finding new and innovative ways to exploit and infiltrate even the most secure systems and the security world is barely able to keep up.
Gone are the simple days of protecting a system from the random individual hacker. Instead, we are now combating a very organized, advanced, and powerful force that comes in many different forms, from script kiddies to nation states, organized crime to hacktivists, and everything in between.
Our job as a security professional is to defend against these organizations and the various techniques they use. Often times this means we need to think like they do look at our system and network from the eyes of an attacker. It also means we need to try to be one step ahead of them at all times. However, this is becoming more and more difficult as the number of internet-connected devices increases and the speed at which people expect new technologies to emerge. See, the faster new technology is developed and created, the less time it has to be thoroughly tested for vulnerabilities, holes, exploits, et cetera. In addition, each new device that connects to a network presents a new point of entry for an attacker that didn't exist before. Even more unsettling are the exploits that haven't even been discovered yet that attackers could using in the wild.
Because of all these variables, the goal for security professionals can never be that of eliminating attacks or breaches, that's impossible.
Now, I'm not saying this means we should just give up. We still need to secure our systems and take all the necessary precautions in order to reduce the threat surface. However, know that if your system is connected to the internet, then it's essentially impossible to protect your network from any and all attacks.
Understanding this, a security professional's goal should instead be to minimize the occurrence of attacks and reduce the damage caused by a breach. In other words, you need to properly secure and protect systems while at the same time understanding that a breach is going to occur. And when it does, you need to be able to identify the moment the breach occurred and stop it as fast as possible.
Doing this requires you take a proactive approach to security. But how do you do that?
Some obvious aspects of this approach include keeping systems up to date, implementing proper policies and procedures, hardening systems and networks, and so on. But another, often times neglected, aspect of this approach includes being informed.
Because the security landscape is ever-changing, you need to be extra diligent in keeping up to date on the most recent vulnerabilities and exploits used by hackers as well as the latest security techniques and technologies used by security

professionals. The internet is an endless supply of information, so be sure to use it. Blogs, news outlets, forums, podcasts, the list goes on, these are all great resources that will help you keep up to date on the latest security trends. Remember, as a security professional, it's your job to try to stay one step ahead of an attacker. You can do this by taking a proactive approach to security. Stay informed, read the landscape, know your systems and network, and understand that you can only protect a network to a point. Beyond that, it's your job to know what an attack looks like and stop it before any substantial damage can occur.

1.1.2 Security Concepts (Lesson Video)

Transcript:

In order to be an effective security professional, you need to be familiar with the concepts and the roles surrounding information security. This will help you understand the industry terms and lingo, and it will also provide a lot of context as you progress through this course.

The first information security concept that you need to be familiar with is that of an asset.
An asset is simply something that has value to an individual or an organization. This can be a physical device, such as a laptop or iPad, or it can be electronic information, such as a pdf document on a server. However, most of the time we're talking about an asset we mean the latter.
For example, let's suppose we have a server in our organization, and on this server there is a database that contains customer information, including credit card numbers and order history. This database has a lot of value to the organization and is therefore considered an asset.
The next security concept that you need to be aware of are threats. Threats represent anything that has the potential to cause the loss of an asset.
And notice I said has the potential to cause the loss of an asset. A threat isn't the actual loss of an asset. It's merely the potentialâ€""the riskâ€""that an asset could be stolen.
A threat can come in many different forms. It can be a virus, a Trojan, an external hacker, an internal employee. Because threats come in all shapes and sizes, sometimes we refer to them as blended threats.
To continue with our example, some threats to our customer database include ransomware, data exfiltrationâ€""which is a fancy way of saying stealing dataâ€""Trojans, and hackers.
Next, we have the threat agent. A threat agent is the actual person or entity that carries out a threat.
When it comes to threat agents, there are a few characteristics, or attributes, that can categorize them. For example, threat agents can be internal or external; they can have a little to no resources or funding, or they can be heavily funded with a lot of resources; they can also be opportunisticâ€""that is, they are simply attacking a system because it has a vulnerabilityâ€""or they can have a specific intent or motive.
Now, within these threat agent categories, there are different types of actorsâ€""the type of entity carrying out the attack. For example, an actor could be an organized crime syndicate trying to steal credit card information. An actor could also be a nation state trying to steal classified information. Even business competitors can be a type of actor who try to steal company secrets in order to gain an economic edge.
One example of a nation state actor you might be familiar with is North Korea. On November 24, 2014, North Korean hackers gained access to Sony Pictures networks and stole confidential information, including employee records, personal emails, and copies of unreleased movies. The information was then released to the public on the internet.
In order for threat agents to carry out a threat, they need an openingâ€""a weakness in the system. This is known as a vulnerability.
For example, a vulnerability could be a disgruntled internal employee who happens to be an information security professional and has an elevated level of access to a server system. Another vulnerability is an enabled USB port.
And the last concept we will talk about is an exploit. An exploit is a procedure, a piece of software, or a sequence of commands that takes advantage of a vulnerability to actually carry out an attack.
For example, say we have an enabled USB port on our customer databaseâ€""first vulnerabilityâ€""and we also have a disgruntled employeeâ€""second vulnerability.
Let's say that the employee decides to use a USB thumb drive to steal the customer database. This is an exploit. The employee used the vulnerability of the enabled USB port and their elevated permissions in order to steal the customer database.
Because security is a constant balancing act between convenience and protection, you will constantly be looking at ways you can mitigate risk and threats while also maintaining an acceptable level of ease of use.

However, by understanding the basic concepts of information security, you will have a much easier time assessing the risks to your systems and identifying the ways in which you can protect it.

1.1.3 Security Introduction Facts

Security is an ongoing process that includes assessing requirements, setting up organizational security systems, hardening and monitoring those systems, responding to attacks in progress, and deterring attackers. If you can summarize the fundamental concepts that underpin security functions, you can contribute more effectively to a security team. You must also be able to explain the importance of compliance factors and best practice frameworks in driving the selection of security controls and how departments, units, and professional roles within different types of organizations implement the security function.

This lesson covers the following topics:

- Security challenges
- Security control types
- Security roles and responsibilities

Security Challenges

In regards to information security, computers, and IT networks, modern-day security challenges include the following:

Challenge	Description
Sophisticated attacks	Sophisticated attacks are complex, making them difficult to detect and thwart. Sophisticated attacks: • Use common internet tools and protocols, making it difficult to distinguish an attack from legitimate traffic. • Vary their behavior, making the same attack appear differently each time.
Proliferation of attack software	A wide variety of attack tools are available on the internet, allowing anyone with a moderate level of technical knowledge to download the tools and run an attack.
Attack scale and velocity	The scale and velocity of an attack can grow to millions of computers in a matter of minutes or days due to its ability to proliferate on the internet. Because modern attacks are not limited to user interactions, such as using a floppy disk to spread an attack from machine to machine, the attacks often affect very large numbers of computers in a relatively short amount of time.

Security Control Types

Information security and cybersecurity assurance are met by implementing security controls. By identifying basic security control types, you will be better prepared to select and implement the most appropriate controls for a given scenario.

All controls are designed to fulfill three main goals: confidentiality, integrity, and availability.

- **Confidentiality** ensures that data is not disclosed to unauthorized persons.
- **Integrity** ensures that data is not modified or tampered with.
- **Availability** ensures the data is available when needed.

Security controls can be classified in different ways to fulfill the goals of your organization.

- Controls can be classified based on the way they are implemented. This includes oversight or managerial controls, operational controls that rely on people, technical or system-based controls, and, finally, non-technical physical controls such as alarms, locks, cameras, etc.
- Another method is to classify controls by the goal or function they perform. This involves controls that prevent attacks before they happen, detect attacks when they occur, and help correct and restore damage caused by attacks.
- Finally, there are controls to cover additional areas such as employee expectations, policies, and employment or disciplinary procedures. Other controls may not actually do anything but are designed to deter malicious actions. There are also times when standards or legal requirements mandate controls that should be used. If it is not possible to implement, an organization can sometimes replace that control with a replacement option that is as good or better than the original.

Security Roles and Responsibilities

You should also be able to describe how specific job roles and organizational structures can implement a comprehensive security program for organizations. IT professionals working in a role with security responsibilities must be competent in a wide range of disciplines, from network and application design to procurement and human resources (HR). The following activities might be typical of such a role:

- Participate in risk assessments and testing of security systems and make recommendations.
- Specify, source, install, and configure secure devices and software.
- Set up and maintain document access control and user privilege profiles.
- Monitor audit logs, review user privileges, and document access controls.
- Manage security-related incident response and reporting.
- Create and test business continuity and disaster recovery plans and procedures.
- Participate in security training and education programs.

The following units are often used to represent the security function within the organizational hierarchy.

A security operations center (SOC) is a location where security professionals monitor and protect critical information assets across other business functions, such as finance, operations, sales/marketing, etc. Because SOCs can be difficult to establish, maintain, and finance, they are usually employed by larger corporations, like a government agency or a healthcare company.

A security operations center (SOC) provides resources and personnel to implement rapid incident detection and response, plus oversight of cybersecurity operations. (Image © gorodenkoff 123RF.com.)

Network operations and cloud computing make ever-increasing use of automation through software code. Traditionally, software code would be the responsibility of a programming or development team. Separate development and operations departments or teams can lead to silos, where each team does not work effectively with the other.

Development and operations (DevOps) is a cultural shift within an organization to encourage much more collaboration between developers and systems administrators. By creating a highly orchestrated environment, IT personnel and developers can build, test, and release software faster and more reliably. DevSecOps extends the boundary to security specialists and personnel, reflecting the principle that security is a primary consideration at every stage of software development and deployment. This is also known as shift left, meaning that security considerations need to be made during requirements and planning phases, not grafted on at the end. The principle of DevSecOps recognizes this and shows that security expertise must be embedded into any development project. Ancillary to this is the recognition that security operations can be conceived of as software development projects. Security tools can be automated through code. Consequently, security operations need to take on developer expertise to improve detection and monitoring.

A dedicated computer incident response team (CIRT)/computer security incident response team (CSIRT)/computer emergency response team (CERT) is a single point of contact for the notification of security incidents. This function might be handled by the SOC, or it might be established as an independent business unit.

1.1.4 Security Job Roles (Lesson Video)

Transcript:

James Stanger

When it comes to the world of cyber security there are a lot of kinda continents available in that world. In other words, there are a lot of different jobs within the realm of cyber security. To tell us more about that we've got Brian Calkin. Brian, how you doin'?
Brian Calkin
I'm doing great, James. How are you?
James Stanger
Doing great, man. Doing great. Brian Calkin is the Chief Technology Officer for CyberWA. And Brian, tell us a bit more about CyberWA yourself. Let's start talking about the world of cyber security jobs that are available to folks.
Brian Calkin
Absolutely, yeah. So I've been in cyber security now for going on 20 years, various roles over the years I've had, to include all the way from a SOC analyst to a person that has built SOCs, run SOCs, led response teams, various management roles,executive roles. And as you mentioned, currently the Chief Technology Officer for a company called CyberWA where we are offering cyber protective services for high-net-worth individuals and celebrities.
James Stanger
So you know when it comes to putting a Security Operations Center together, more importantly you have to have workers at various tiers, various levels. Is a Security Operations Center analyst, a SOC analyst, a good first cyber security job?
Brian Calkin
I'd actually argue it's the best first cyber security job, but I can stand by that pretty strongly having been one, and having hired many and managed many of them over the years. I think that as folks are getting started in cyber security, and they may be not exactly sure where their passions might lie, whether they want to dive into penetration testing, or vulnerability assessments, or incident response, I think that you get enough exposure to each of those various individual sort of work items with being a SOC analyst. You just get this broad experience that you can apply later on in your career.
James Stanger
I've known quite a few beginning SOC analysts. They basically take tickets, as it were, oftentimes, right? And then they determine if they can handle it themselves, if it really is a big thing, and then escalate it. Is that, in general terms, what you're gonna do with that particular job role?

Brian Calkin
Absolutely, yeah. So generally the way it works is we will hire somebody, you know, either straight out of college, or maybe new to cyber security in general, as what we call an entry level or a tier one analyst. So these are the folks that are
getting the broadest experience in handling most of the, if not all of the inbound requests from customers, whether they be internal customers to your organization, or, you know, external customers if you're running an external shop. But yeah, fielding emails, fielding phone calls, and running down things like, you know, requests for analysis or requests for assistance. And then yeah, if they get into a bind or a tough spot where they've sort of hit the limits of what they know and can do, then they would escalate up to the tier two analyst.
James Stanger
You know, there are other job roles out there, you know, a pen tester for example, or security administrator. Let's talk about a couple of those real quick.
Brian Calkin
Yeah, sure. Those are always popular too, and folks, they like the pen testing roles, security admin roles, security engineer roles. So for instance a penetration tester would be somebody that is identifying vulnerabilities in a system, or a network, or a networkperimeter, and then trying to put on the role of the hacker or the black hat, and then try to exploit those vulnerabilities to try to get in, you know. And the purpose, of course, is to let the organization know they found this vulnerability and they were able to exploit it, versus a security engineer would be, so if a pen tester is playing offense, the engineer would be playing defense. The security engineer is preparing those systems to be as secure as possible as they're being deployed. Think, you know, deploying a new server and hardening it for best practice or per recommendations in order to make it as resilient as possible to a cyber attack.
James Stanger
Additional job roles can include things like, you know, a help desk analyst, or a vulnerability manager, things like that. But tell me a bit more about some of those.
Brian Calkin
Yeah, sure. Vulnerability assessment, vulnerability manager is another really popular one that, you know, I've worked with a number of those folks over the years. So these are the folks that are essentially doing things like, or could include running vulnerability scans, looking at various vulnerabilities that an organization might have within their environment, and then assessing, you know, what is the impact should this vulnerability get exploited, making sure the vulnerability can be patched, if it maybe can't be patched, and what else they could do to help mitigate the vulnerability. You know, is there some other mitigating factor in play they could apply to kind of close that hole, you know, if you will.
James Stanger
Brian, thank you so much for talking about just a few of the many job roles that are available within cyber security. Sure appreciate it, man.
Brian Calkin
Yeah, you're welcome. Thanks, James.

1.1.5 Practice Questions (Section Quiz)

q_sec_intro_cia_secp8

Which of the following are often identified as the three main goals of security? (Select three.)

Answers:

- ***Integrity**
- ***Confidentiality**
- Employees
- ***Availability**
- Policies
- Assets
- Non-repudiation

Explanation:

The acronym CIA refers to confidentiality, integrity, and availability in respect to security. These are often identified as the three main goals of any security-oriented task.

Non-repudiation provides validation of a message's origin.

Policies are the rules an organization implements to protect information.

Employees can be the most overlooked, yet most dangerous, threat agent because they have greater access to information assets than anyone on the outside trying to break in.

An asset is something that has value to a person or organization, such as sensitive information in a database.

q_sec_intro_cirt_secp8

A large multinational corporation has recently experienced a significant data breach. The breach was detected by an external cybersecurity firm, and the corporation's IT department was unable to prevent or detect the breach in its early stages.

The CEO wants to ensure that such a breach does not happen again and is considering several options to enhance the company's security posture.

Which of the following options would be the MOST effective in preventing and detecting future data breaches?

Answers:

- Hiring an external cybersecurity firm to conduct regular penetration testing.
- ***Implementing a dedicated Computer Incident Response Team (CIRT).**
- Increasing the budget for the IT department to purchase more advanced security software.
- Conducting regular cybersecurity training for all employees.

Explanation:

Implementing a dedicated Computer Incident Response Team (CIRT) is the most effective option. A CIRT is a group of experts who respond to security incidents. They have the skills and authority to respond quickly and effectively to security incidents, minimizing the impact and preventing further damage. This team would also be responsible for monitoring the company's network for signs of a breach, allowing them to detect and respond to incidents more quickly in the future.

Hiring an external cybersecurity firm to conduct regular penetration testing is a good practice, but it is not sufficient on its own. Penetration testing can identify vulnerabilities, but it does not provide continuous monitoring or incident response capabilities.

Increasing the budget for the IT department to purchase more advanced security software can be helpful, but software alone cannot fully protect a company from data breaches. Effective security requires a combination of technology, processes, and people.

Conducting regular cybersecurity training for all employees is important, but it is not enough to prevent and detect data breaches. Employees can help prevent breaches by following good security practices, but they cannot be expected to have the specialized skills needed to detect and respond to sophisticated attacks.

q_sec_intro_confidentiality_secp8

A user copies files from her desktop computer to a USB flash device and puts the device into her pocket.

Which of the following security risks is MOST pressing?

Answers:

- ***Confidentiality**
- Availability
- Non-repudiation
- Integrity

Explanation:

Confidentiality ensures that data is not disclosed to unintended persons. Removable media poses a big threat to confidentiality because it makes it easy to remove data and share it with unauthorized users.

Availability ensures that data is available when it is needed. Copying files to a server that includes malware could threaten the data's availability if the malware deletes or corrupts the data.

Integrity ensures that data is not modified or tampered with.

Non-repudiation provides validation of a message's origin.

q_sec_intro_devops_secp8

You are the Chief Information Security Officer (CISO) at a tech company. Your company is facing issues with silos between the development and operations teams, leading to inefficiencies and security vulnerabilities.

Which approach should you adopt to encourage collaboration and integrate security considerations at every stage of software development and deployment?

Answers:

- Implementing a new security policy
- Establishing a Security Operations Center (SOC)
- ***Adopting a Development and Operations (DevOps) approach**
- Outsourcing security to a third-party vendor

Explanation:

Adopting a Development and Operations (DevOps) approach is the correct answer. The DevOps approach encourages collaboration between developers and systems administrators, integrating security considerations at every stage of software development and deployment. This can help break down silos and improve both efficiency and security.

While implementing a new security policy can help improve security, it does not directly address the issue of silos between the development and operations teams.

A Security Operations Center (SOC) is a dedicated team that monitors and protects critical information assets across the organization. However, it does not directly address the issue of silos between the development and operations teams.

While outsourcing can be a viable option for some organizations, it does not directly address the issue of silos between the development and operations teams. Furthermore, it may not provide the same level of control and oversight as an integrated internal team.

q_sec_intro_integrity_secp8

Your computer system is a participant in an asymmetric cryptography system. You've created a message to send to another user. Before transmission, you hash the message and encrypt the hash using your private key. You then attach this encrypted hash to your message as a digital signature before sending it to the other user.

In this example, which protection does the hashing activity provide?

Answers:

- Confidentiality
- ***Integrity**
- Availability
- Non-repudiation

Explanation:

Hashing of any sort, including within a digital signature, provides data integrity.

Signing the message with the private key creates non-repudiation.

A digital signature activity, as a whole, does not provide protection for confidentiality because the original message is sent in cleartext.

No form of cryptography provides protection for availability.

q_sec_intro_managerial_controls_secp8

Which of the following is a method of implementing security controls?

Answers:

- ***Managerial controls**
- Financial controls
- Marketing controls
- Sales controls

Explanation:

Managerial controls is the correct answer. Managerial controls are administrative actions that define the organization's security posture through policy, guidelines, standards, procedures, and other forms of documentation. They provide a framework for operation and are essential for the organization's security program.

While financial controls are important in a business context, they are not a method of implementing security controls. Financial controls manage the accuracy and reliability of financial transactions and reporting, not the security of an organization's information systems.

Marketing controls are not a method of implementing security controls. They are methods used by marketing departments to analyze and measure their strategies and campaigns, not to secure an organization's information systems.

Sales controls are not a method of implementing security controls. They are methods used by sales departments to manage and track sales activities, not to secure an organization's information systems.

q_sec_intro_scale_and_velocity_secp8

Which of the following security challenges refers to the rapid and broad spread of an attack, often affecting a large number of computers in a relatively short amount of time?

Answers:

- Sophisticated attacks
- Proliferation of attack software
- ***Attack scale and velocity**
- Data encryption

Explanation:

Attack scale and velocity is the correct answer. It refers to the scale and speed of an attack, which can grow to millions of computers in a matter of minutes or days due to its ability to proliferate on the internet.

Sophisticated attacks are complex attacks that are difficult to detect and thwart. They use common internet tools and protocols, making it difficult to distinguish an attack from legitimate traffic. They also vary their behavior, making the same attack appear differently each time. However, this option does not specifically refer to the speed and scale of an attack.

Proliferation of attack software refers to the wide variety of attack tools available on the internet, allowing anyone with a moderate level of technical knowledge to download the tools and run an attack. While this can contribute to the scale and velocity of an attack, it is not the same thing.

While data encryption is a crucial aspect of cybersecurity, it is not a security challenge. Instead, it is a method used to protect data from unauthorized access.

q_sec_intro_soc_secp8

You are the Chief Information Security Officer (CISO) at a large corporation. Your company is expanding rapidly and the complexity of managing security across different business functions is increasing.

You need a dedicated team to monitor and protect critical information assets across the organization.

Which of the following would be the MOST effective solution?

Answers:

- Implementing a new security policy
- ***Establishing a Security Operations Center (SOC)**
- Hiring more IT staff
- Outsourcing security to a third-party vendor

Explanation:

Establishing a Security Operations Center (SOC) is the correct answer. A SOC is a dedicated team that monitors and protects critical information assets across the organization. It provides resources and personnel to implement rapid incident detection and response, plus oversight of cybersecurity operations.

While implementing a new security policy can help improve security, it does not provide the dedicated monitoring and protection needed for critical information assets across different business functions.

While hiring more IT staff can help manage the increased workload, it does not provide the dedicated monitoring and protection needed for critical information assets across different business functions.

While outsourcing can be a viable option for some organizations, it may not provide the same level of control and oversight as a dedicated internal team like a SOC.

q_sec_intro_sophisticated_attacks_secp8

You are the head of the cybersecurity team at a large corporation. You notice an increase in network traffic that appears to be legitimate but is causing a slowdown in your systems.

Upon further inspection, you find that the traffic patterns vary each time, making it difficult to distinguish from normal traffic.

What type of security challenge are you MOST likely facing?

Answers:

- Proliferation of attack software
- Attack scale and velocity
- ***Sophisticated attack**
- Data breach

Explanation:

A sophisticated attack is the correct answer. Sophisticated attacks are complex and often use common internet tools and protocols, making them difficult to distinguish from legitimate traffic. They also vary their behavior, making the same attack appear differently each time, which aligns with the observed traffic patterns.

Proliferation of attack software refers to the wide availability of attack tools on the internet, which allows anyone with moderate technical knowledge to download the tools and run an attack. While this could potentially cause an increase in network traffic, it does not account for the varying traffic patterns observed.

Attack scale and velocity refers to the rapid and widespread proliferation of an attack, often affecting a large number of computers in a short amount of time. While this could cause a slowdown in systems, it does not account for the varying traffic patterns observed.

A data breach refers to an incident where information is accessed without authorization. While a data breach could potentially result from the security challenge faced, it does not in itself explain the varying traffic patterns observed.

q_sec_intro_technical_control_secp8

You are the Chief Information Security Officer (CISO) at a large corporation. You have been tasked with implementing a new security control to protect sensitive customer data.

The control must be able to automatically detect and prevent unauthorized access to the data.

Which type of control should you implement?

Answers:

- Managerial control
- Operational control

- ***Technical control**
- Physical control

Explanation:

Technical control is the correct answer. Technical controls involve the use of technology to reduce vulnerabilities. This includes controls that can automatically detect and prevent unauthorized access, such as intrusion detection systems, firewalls, and access control mechanisms.

Managerial controls are administrative actions that define the organization's security posture through policy, guidelines, standards, procedures, and other forms of documentation. While they are crucial for setting the overall direction of an organization's security, they do not automatically detect and prevent unauthorized access to data.

Operational controls are the day-to-day procedures and mechanisms that protect an organization's assets. While they can include procedures for detecting and preventing unauthorized access, they do not typically do this automatically.

Physical controls involve measures to prevent physical access to assets, such as locks, fences, and security guards. While they are crucial for protecting physical assets, they do not automatically detect and prevent unauthorized access to digital data.

1.2 Security Controls

As you study this section, answer the following questions:

- What are the three types of control categories?
- What are preventative controls?
- What is the difference between corrective and compensating controls?

The key terms for this section include:

Term	Definition
Security control	A technology or procedure put in place to mitigate vulnerabilities and risk and to ensure the confidentiality, integrity, and availability (CIA) of information.
Managerial	A category of security control that provides oversight of information systems.
Operational	A category of security control that is implemented by people.
Technical	A category of security control that is implemented as a system.
Physical	A category of security control that is implemented by hardware used to deter or detect, such as as alarms, gateways, locks, lighting, and security cameras.
Preventive	A type of security control that acts before an incident to eliminate or reduce the likelihood that an attack can succeed.

Access control lists (ACLs)	The collection of access control entries (ACEs) that determines which subjects (user accounts, host IP addresses, and so on) are allowed or denied access to the object and the privileges given (read-only, read/write, and so on).
Detective	A type of security control that acts during an incident to identify or record that it is happening.
Corrective	A type of security control that acts after an incident to eliminate or minimize its impact.
Directive	A type of control that enforces a rule of behavior through a policy or contract.
Deterrent	A type of security control that discourages intrusion attempts.
Compensating	A security measure that takes on risk mitigation when a primary control fails or cannot completely meet expectations.
Chief Information Officer (CIO)	A company officer with the primary responsibility of managing information technology assets and procedures.
Chief Technology Officer (CTO)	A company officer with the primary role of making effective use of new and emerging computing platforms and innovations.
Chief Security Officer (CSO)	Typically, the job title of the person with overall responsibility for information assurance and systems security.
Information Systems Security Officer (ISSO)	Organizational role with technical responsibilities for implementation of security policies, frameworks, and controls.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.1 Compare and contrast various types of security controls.
	• Categories
	  ○ Technical
	  ○ Managerial
	  ○ Operational
	  ○ Physical
	• Control types
	  ○ Preventive
	  ○ Deterrent

	○ Detective ○ Corrective ○ Compensating ○ Directive

1.2.1 Control Categories and Types (Lesson Video)

Transcript:

In this video, I'll discuss security control categories and types. As a security professional, your job is to protect the company's assets and business. This means that you're there to prevent hazardous events or to minimize their damage. Let's get started.

Managerial controls consist of management techniques and administrative procedures. These can include security policies, hiring policies, disaster recovery plans, or DPRs, and business continuity plans, or BCPs. You should always include written plans for how and when to escalate events and when you need to consult outside help.
Operational controls are ones that the security team performs everyday. These might include reviewing network monitoring data, ensuring that security cameras are working, and requiring visitors to sign in.
Technical controls are based around software, applications, and security appliances. These controls vary greatly by company. Some examples you might see in your line of work are IDSs, IPSs, access control apps, and ASAs.
We use preventative controls for security breaches, but you might see these tools used with other control types as well. The easiest preventative control is an Advance Network Appliance, which is sometimes also called an Adaptive Security Appliance, or ASA. This is basically a firewall and router combination that's capable of hosting IDSs and IPSs. A less expensive preventative control is a simple updated antivirus, considering that these have improved drastically over the years. Office access control is also an excellent preventative control, especially if you're using biometrics.
Detective controls inform the security team of an event that's occurring or provide them with logs and artifacts to them help investigate the event further. Examples would be network monitoring applications, log collectors, real-time monitoring alerts, and intrusion detection systems, or IDSs.
Corrective controls are those that attempt to fix any damage caused by an event. These tools work during and after the course of the event. Think of this as a form of risk mitigation. For example, an intrusion prevention system, or IPS, is designed to intercept packets that are potentially malicious and either drop or isolate them. Another example is endpoint protection, which works to stop malicious data by looking at its signature or behavior.
The deterrent control type discourages malicious actors from trying to breach a network. The more deterrents you have, the less likely it is that anyone will try. These could include internal security
policies, access-protected doors for a server room, entry-point access restriction, biometric sensors, man traps, security cameras, security training, and security guards. Remember, the stronger the deterrents, the less likely it is that a breach occurs.
Physical deterrents keep unauthorized people from physically accessing a company's assets. So locked doors, proximity cards, fences, cameras, and guards are all ways to physically protect your network. Motion detectors for after-hours monitoring is another example.
Please note that device management is often overlooked by companies. With the prevalence of mobile devices, your company should be able to remotely wipe any devices that are lost or stolen. Let's look at some scenarios.
Let's say that a biometric thumbprint scanner has been installed in the lobby of your building. It requires everyone to scan their thumb to gain access. A malicious actor tries to enter, but their print doesn't match. An access-denied alert is flashed on the screen, and the IT team is fully aware of what's happening. Usually, the imposter doesn't stay around to be questioned!
Another common ploy is for an imposter to try and take advantage of a person's kindness. In this case, an imposter tries to gain entrance from a legitimate employee's proximity badge by claiming to have forgotten theirs. Hopefully, the employee remembers their security training and politely tells the imposter, "No."
That's it for this lesson. In this lesson, we learned about control categories as well as control types. Categories include managerial, operational, and technical. The security control types we learned about are preventative, corrective, deterrent, and physical. You should use all of these tools together to help you create a secure network.

1.2.2 Control Categories and Types Facts

Information security and cybersecurity assurance are met by implementing security controls. By identifying basic security control types, you will be better prepared to select and implement the most appropriate controls for a given scenario. You should also be able to describe how specific job roles and organizational structures can implement a comprehensive security program for organizations.

This lesson covers the following topics:

- Security control categories
- Security control functional types

Security Control Categories

Information and cybersecurity assurance usually takes place within an overall process of business risk management. Implementation of cybersecurity functions is often the responsibility of the IT department. There are many ways of thinking about how IT services should be governed to fulfill overall business needs. Some organizations have developed IT service frameworks to provide best practice guides for implementing IT and cybersecurity. These frameworks can shape company policies and provide checklists of procedures, activities, and technologies that represent best practices. Collectively, these procedures, activities, and tools can be referred to as security controls.

A security control is designed to give a system or data asset the properties of confidentiality, integrity, availability, and non-repudiation. Controls can be divided into four broad categories based on the way the control is implemented:

- **Managerial** — the control gives oversight of the information system. Examples could include risk identification or a tool allowing the evaluation and selection of other security controls.
- **Operational** — the control is implemented primarily by people. For example, security guards and training programs are operational controls.
- **Technical** — the control is implemented as a system (hardware, software, or firmware). For example, firewalls, antivirus software, and OS access control models are technical controls.
- **Physical** — controls such as alarms, gateways, locks, lighting, and security cameras that deter and detect access to premises and hardware are often placed in a separate category from technical controls.

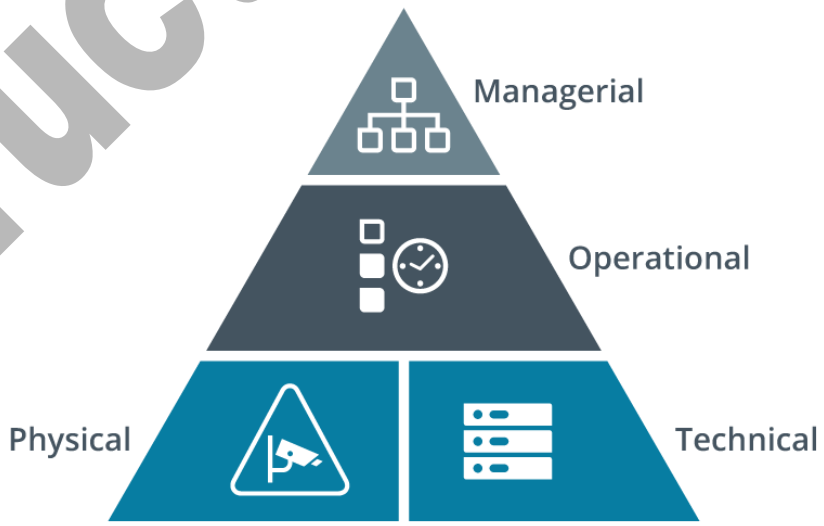

Categories of security controls

Although it uses a different scheme, be aware of how the National Institute of Standards and Technology (NIST) classifies security controls (csrc.nist.gov/publications/detail/sp/800-53/rev-5/final).

For example, as of NIST 800-53 rev 4, the class designations of technical, operational, and managerial were removed from the control families list. Instead, they were redefined as properties of individual controls within a family. They are included to help familiarize learners with the basic concepts presented in 800-53 and due to the continued use of this terminology by many organizations and publications. Be aware that terminology usage and practice are always evolving.

Security Control Functional Types

As well as a category, a security control can be defined according to the goal or function it performs:

- **Preventive** — the control acts to eliminate or reduce the likelihood that an attack can succeed. A preventive control operates before an attack can take place. Access control lists (ACLs) configured on firewalls and file system objects are preventive-type technical controls. Antimalware software acts as a preventive control by blocking malicious processes from executing.
- **Detective** — the control may not prevent or deter access, but will identify and record an attempted or successful intrusion. A detective control operates during an attack. Logs provide one of the best examples of detective-type controls.
- **Corrective** — the control eliminates or reduces the impact of a security policy violation. A corrective control is used after an attack. A good example is a backup system that restores data damaged during an intrusion. Another example is a patch management system that eliminates the vulnerability exploited during the attack.

While most controls can be classed functionally as preventive, detective, or corrective, a few other types can be used to define other cases:

- **Directive** — the control enforces a rule of behavior, such as a policy, best practice standard, or standard operating procedure (SOP). For example, an employee's contract will set out disciplinary procedures or causes for dismissal if they do not comply with policies and procedures. Training and awareness programs can also be considered as directive controls.
- **Deterrent** — the control may not physically or logically prevent access, but it psychologically discourages an attacker from attempting an intrusion. This could include signs and warnings of legal penalties against trespass or intrusion.
- **Compensating** — the control is a substitute for a principal control, as recommended by a security standard. It affords the same (or better) level of protection but uses a different methodology or technology.

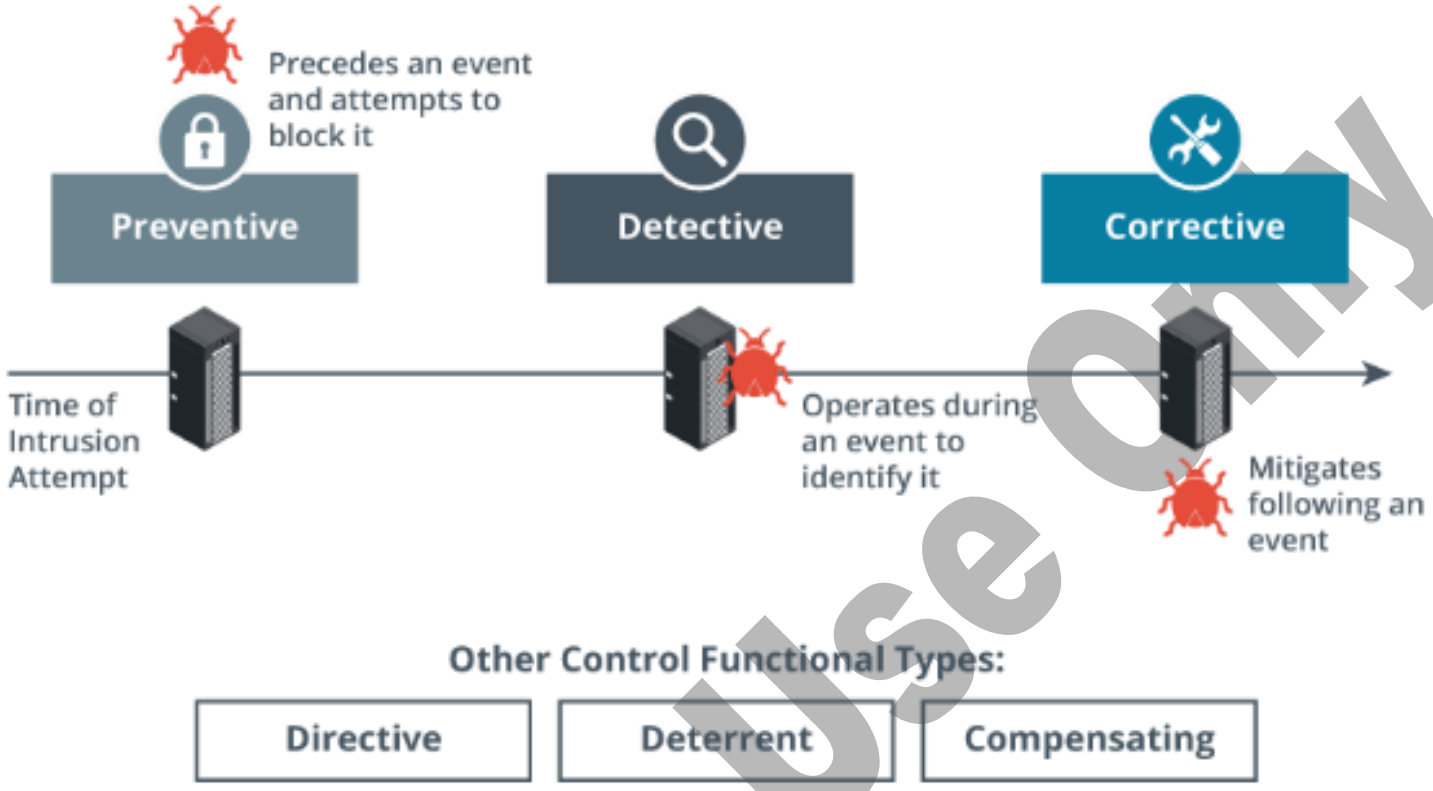

Functional types of security controls. (Images © 123RF.com.)

A security policy is a formalized statement that defines how security will be implemented within an organization. It describes the means the organization will take to protect the confidentiality, availability, and integrity of sensitive data and resources.

The implementation of a security policy to support the goals of the CIA triad might be very different for a school, a multinational accountancy firm, or a machine tool manufacturer. However, each of these organizations, or any other organization (in any sector of the economy, whether profit-making or non-profit-making), should have the same interest in ensuring that its employees, equipment, and data are secure against attack or damage. An organization that develops security policies and uses framework-based security controls has a strong security posture.

As part of the process of adopting an effective organizational security posture, employees must be aware of their responsibilities. The structure of security responsibilities will depend on the size and hierarchy of an organization, but these roles are typical.

- Overall responsibility for the IT function lies with a Chief Information Officer (CIO) . This role might also have direct responsibility for security. Some organizations will also appoint a Chief Technology Officer (CTO) , with more specific responsibility for ensuring the effective use of new and emerging IT products and solutions to achieve business goals.
- In larger organizations, internal responsibility for security might be allocated to a dedicated department run by a Chief Security Officer (CSO) or Chief Information Security Officer (CISO).
- Managers may have responsibility for a domain, such as building control, web services, or accounting.
- Technical and specialist staff have responsibility for implementing, maintaining, and monitoring the policy. Security might be made of a core competency of systems and network administrators, or there may be dedicated security administrators. One such job title is Information Systems Security Officer (ISSO) .
- Nontechnical staff have the responsibility of complying with policy and with any relevant legislation.
- External responsibility for security (due care or liability) lies mainly with directors or owners, though again, it is important to note that all employees share some measure of responsibility.

NIST's National Initiative for Cybersecurity Education (NICE) categorizes job tasks and job roles within the cybersecurity industry (gov/itl/applied-cybersecurity/nice/nice-framework-resource-center).

1.2.3 Security Control and Framework Types

1.2.4 Practice Questions (Section Quiz)

q_ctrl_cattypes_categories_secp8

Which of the following are control categories? (Select three.)

Answers:

- ***Managerial**
- Preventative
- ***Operational**
- Deterrent
- ***Technical**
- Compensating
- Physical

Explanation:

Control categories are:

- Managerial controls (consist of management techniques and administrative procedures)
- Operational controls (performed everyday by the security team)
- Technical controls (based around software, applications, and security appliances)

The remaining items are all control types. Control types consist of different strategies to prevent, detect, mitigate, and correct any network breach.

q_ctrl_cattypes_compensating_01_secp8

Which of the following BEST describes compensating controls?

Answers:

- ***Partial control solution that is implemented when a control cannot fully meet a requirement.**
- Monitors network activity and informs the security team of a potential security event.
- Attempts to fix any controls that aren't working properly.
- Discourages malicious actors from attempting to breach a network.

Explanation:

Compensating controls are a partial control solution that is implemented when a control cannot fully meet a requirement.

Detective controls monitor network activity and inform the security team of a potential security event.

Corrective controls attempt to fix any controls that aren't working properly.

Deterrent controls discourage malicious actors from attempting to breach a network.

q_ctrl_cattypes_compensating_02_secp8

The IT director at a financial institution focuses on implementing compensating managerial controls to augment the institution's existing security framework.

If a mandated control cannot be put into place, which of the following compensating controls should an analyst recommend as a sufficient substitute?

Answers:

- ***Isolating a critical system that cannot be patched.**
- Regular employee training on cybersecurity best practices.
- Using biometric access controls on all company systems.
- An automated system that scans and patches software vulnerabilities.

Explanation:

In a scenario where the analyst cannot patch a critical system due to operational reasons, isolating it from the network can serve as a compensating control. This offers an equivalent level of protection by ensuring potential cyber threats cannot reach the system.

While employee training is an important part of any organization's security controls, it classifies as an operational control more than a managerial one.

Biometric access controls are technical controls, not managerial ones. They use hardware and software to restrict access to systems and data.

While an automated system for scanning and patching software vulnerabilities is a critical part of any organization's cybersecurity strategy, it is a technical control, not a managerial control.

q_ctrl_cattypes_compensating_03_secp8

After finding some of the company's confidential data on the internet, a software team is drafting a policy on vulnerability response and remediation.

What remediation practice refers to measures put in place to mitigate the risk of a vulnerability when the team cannot directly eliminate it?

Answers:

- ***Compensating controls**
- Patching
- Insurance
- Segmentation

Explanation:

Compensating controls refer to measures put in place to mitigate the risk of a vulnerability when security teams cannot directly eliminate it or when direct remediation is not immediately possible.

Patching is one of the most straightforward and effective remediation practices. It involves applying updates and patches to software or systems to fix known vulnerabilities.

Insurance provides financial protection in case of a security breach resulting from a vulnerability and is important in a comprehensive risk management strategy, complementing technical controls with financial risk transfer.

Segmentation involves dividing a network into separate pieces to contain potential security breaches.

q_ctrl_cattypes_compensating_04_secp8

An organization changes its security posture after a breach and wants to enhance encryption by putting meaures in place to mitigate risk exposures that cannot be directly eliminated by the cyber security team.

What type of control is being observed in this situation?

Answers:

- ***Compensating**
- Technical
- Administrative
- Detective

Explanation:

Compensating controls refer to measures put in place to mitigate the risk of a vulnerability when security teams cannot directly eliminate it or when direct remediation is not immediately possible, such as additional monitoring or enhanced encryption.

The security manager would implement a technical control in operating systems, software, and security appliances, with examples including an access control list (ACL) or an intrusion prevention system (IPS).

The security manager uses administrative security controls to dictate behavior through guidelines, policies, and procedures.

While a detective control may not deter or prevent access, it records and identifies any successful or attempted intrusion, exemplified by a security camera system.

q_ctrl_cattypes_corrective_01_secp8

A company moved its office supplies to another room and instituted a new security system for entry. The company implemented this after a recent server outage.

What category of security control BEST describes the function of this recent implementation?

Answers:

- ***Corrective**
- Preventive
- Detective
- Operational

Explanation:

Corrective controls eliminate or reduce the impact of a security policy violation. A corrective control occurs after an attack. For this scenario, segregating server space access from common access would be corrective.

Preventive controls eliminate or reduce the likelihood that an attack can succeed. The company implements this control to avert a potential incident from occurring.

Detective controls may not prevent or deter access, but they will identify and record an attempted or successful intrusion. A security camera would be a type of detective control.

Operational controls involve people, such as hiring security guards and performing training programs.

q_ctrl_cattypes_corrective_02_secp8

After encountering a cyber attack, an organization uses a monitoring solution that automatically restarts services after it has detected the system has crashed.

What type of functional security control is the company implementing?

Answers:

- ***Corrective**
- Compensating
- Technical
- Managerial

Explanation:

Corrective control actively responds to an incident, fixes it, and prevents it from happening again. Antivirus software exemplifies a corrective control.

Rather than preventing an attack, a compensating control actively restores the functionality of systems via alternative methods, such as using a backup.

The implementation of technical controls occurs in operating systems, software, and security appliances such as access control lists (ACLs) or intrusion prevention systems (IPS).

Managerial security controls actively shape behavior through policies, procedures, and guidance. A fair use policy serves as an example of such a control.

q_ctrl_cattypes_detective_secp8

The chief security officer (CSO) at a financial organization wants to implement additional detective security controls.

Which of the following would BEST represent this type of control?

Answers:

- ***Installation of surveillance camera.**
- Enforcement of access control mechanisms.
- Implementation of biometric authentication systems.
- Performing regular system backups.

Explanation:

Surveillance cameras are physical controls that act as detective mechanisms, helping to identify unauthorized access or activities within the monitored area.

While it is an important security measure, enforcing access control mechanisms primarily serves as a preventive control. For example, it prevents unauthorized physical access but does not detect issues or breaches actively.

Although a biometric authentication system is an important control, it primarily serves as a preventive control. Biometric authentication prevents unauthorized access by ensuring verified individuals can access certain areas or data, but it does not actively detect breaches or security incidents.

Regular system backups are generally a corrective control that does not detect issues. Instead, they provide a means to recover after a security incident.

q_ctrl_cattypes_deterrent_01_secp8

Which type of control is used to discourage malicious actors from attempting to breach a network?

Answers:

- Detective
- Physical
- Preventative
- ***Deterrent**

Explanation:

The deterrent control type discourages malicious actors from trying to breach a network. The more deterrents are implemented, the less likely it is that anyone tries. These could include internal security policies, access-protected doors for a server room, entry-point access restriction, biometric sensors, man traps, security cameras, security training, and security guards.

Detective controls monitor network activity and inform the security team of a potential security event. Detective controls also log activities and provide artifacts to help investigate the event. Intrusion detection systems are an example of detective controls.

Physical deterrents keep unauthorized people from physically accessing a company's assets. Locked doors, proximity cards, fences, cameras, and guards are all ways to physically protect a network.

Preventative controls, such as an IPS, are used to prevent security breaches.

q_ctrl_cattypes_deterrent_02_secp8

A manufacturing company is looking to enhance its security measures by implementing deterrent controls in its facility, specifically the server room.

Which of the following options would be MOST effective?

Answers:

- ***Placing visible signs indicating surveillance and severe penalties for unauthorized entry.**

- Installing a network-connected smoke detector system in the server room.
- Applying reflective window film to the server room windows.
- Introducing a key control system for office desks.

Explanation:

Placing visible signs indicating surveillance and severe penalties for unauthorized entry is an example of a physical control. This kind of deterrence can discourage unauthorized physical access by setting clear consequences.

Installing a network-connected smoke detector system in the server room is a physical control, but it's designed to detect fire hazards, not deter unauthorized access.

Applying reflective window film to the server room windows is a physical control aimed at enhancing privacy, but it doesn't deter unauthorized access once inside the building.

Introducing a key control system for office desks is a physical control, but it's more relevant for securing individual workstations and personal belongings than deterring unauthorized access to a server room.

q_ctrl_cattypes_deterrent_03_secp8

A user in a company wants a new USB flash drive. Rather than requesting one through the proper channel, the user obtains one from one of the company's storage closets.

Upon approaching the closet door, the user notices a warning sign indicating cameras are in use.

What is the control objective of the observed sign?

Answers:

- ***Deterrent**
- Preventive
- Detective
- Corrective

Explanation:

A deterrent control may not physically or logically prevent access, but psychologically, it discourages an attacker from attempting an intrusion. A warning sign is an example of a deterrent control.

A preventive control physically or logically restricts unauthorized access. A system password and physical door lock are examples of preventive controls.

A detective control may not prevent or deter access, but it will identify and record any attempted or successful intrusion. A security camera system is an example of a detective control.

A corrective control responds to and fixes an incident. It may also prevent the reoccurrence of the incident. Antivirus software is an example of a corrective control.

q_ctrl_cattypes_directive_01_secp8

The head of IT security at a financial institution is working to enhance the directive controls in place within the company.

Which of the following should the institution implement?

Answers:

- ***Building access procedures**
- Biometric access control systems
- Intrusion detection systems (IDS)
- Closed-circuit television surveillance cameras

Explanation:

Building access procedures serve as an example of directive controls. They provide guidance and direction on the steps for gaining access to the building, helping to achieve the desired security outcome.

Biometric access control systems primarily function as technical and physical controls and do not serve as directive controls. They control access to physical or digital resources using biological characteristics like fingerprints or retinal scans.

Intrusion detection systems typically fall under the detective control category. While they can identify potentially harmful activities within a network, they do not offer guidance or direction.

Closed-circuit television surveillance cameras, which function primarily as physical controls, monitor and detect unauthorized access or activity. Although they can act as deterrents, they do not get classified as directive controls.

q_ctrl_cattypes_directive_02_secp8

The security operations manager of a multinational corporation focuses on enhancing directive operational controls.

Which of the following should the manager implement?

Answers:

- ***User awareness and training programs.**
- Firewall to block unauthorized network traffic.
- Surveillance cameras installed around the premises.
- Regular vulnerability assessments using automated tools.

Explanation:

User awareness and training programs are examples of directive operational controls. These controls are about guiding behavior towards specific outcomes, such as improving security awareness and adherence to protocols.

A firewall is a preventive control that serves to block unauthorized network traffic. Despite its importance, it doesn't guide behavior and therefore is not a directive control.

Surveillance cameras serve to monitor and detect potential security incidents. While they contribute to overall security, they do not guide behavior and are not directive controls.

Regular vulnerability assessments help in identifying potential security weaknesses. Although critical, they do not guide behavior, thus they are not considered directive controls.

q_ctrl_cattypes_directive_03_secp8

An information technology manager conducted an audit of the company's support tickets. The manager noticed a trend with the tickets, where the majority were for new computer setups.

What security control function would the manager's implementation of a new standard operating procedure have?

Answers:

- Compensating
- Deterrent
- ***Directive**
- Corrective

Explanation:

A directive control enforces a rule of behavior, such as a policy, best practice standard, or standard operating procedure (SOP).

Compensating controls are a substitute for a principal control, as recommended by a security standard, and afford the same (or better) level of protection. However, they use a different methodology or technology.

Deterrent controls may not physically or logically prevent access, but it psychologically discourages an attacker from attempting an intrusion. Deterrent controls could include signs and warnings of legal penalties against trespass or intrusion.

Corrective controls eliminate or reduce the impact of a security policy violation. A corrective control is used after an attack. A good example is a backup system that restores data that was damaged during an intrusion.

q_ctrl_cattypes_managerial_secp8

Which type of control makes use of policies, DPRs, and BCPs?

Answers:

- Operational
- Technical
- ***Managerial**
- Preventative

Explanation:

Managerial controls consist of management techniques and administrative procedures. These can include security policies, hiring policies, disaster recovery plans (DPRs), and business continuity plans (BCPs).

Operational controls are ones that the security team performs daily.

Technical controls are based around software, applications, and security appliances.

Preventative controls, such as an IPS, are used to prevent security breaches.

q_ctrl_cattypes_operational_secp8

An acceptable use policy requires the system to encrypt confidential information while in transit. All employees must use secure email when exchanging proprietary information with external vendors.

Which of the following describes this type of acceptable use policy?

Answers:

- ***Operational**
- Managerial
- Technical
- Preventive

Explanation:

Operational controls like this acceptable use policy focus on procedures and responsibilities that are well defined and executed by people. They help to ensure the security of an organization's day-to-day operations.

Managerial controls establish strategies, goals, and objectives for an organization's overall security program. They often include risk assessment and the review of security controls.

Technical controls involve the use of technology to control user access and protect information systems. They include network and system monitoring, firewalls, and intrusion detection systems.

Preventive controls are proactive measures designed to stop potential security incidents from happening. Examples include firewall configurations, user access controls, and security awareness training.

q_ctrl_cattypes_physical_01_secp8

When setting up a new server room for sensitive data storage, a tech company seeks to enhance preventive measures against unauthorized access.

Which measure would be MOST effective for this purpose?

Answers:

- ***Physical security**
- Video surveillance
- Server encryption
- Intrusion detection system (IDS)

Explanation:

Physical security is the first line of defense against unauthorized access to a server room. Physical security can include locked doors, access control systems, and security guards.

Although video surveillance offers a potent deterrent, its principal function is reactive rather than preventive. However, video surveillance does not preemptively stave off unauthorized access.

Server encryption primarily pertains to digital breaches and does not extend its protective purview to inhibit physical trespassing into a server room.

Although an IDS is invaluable for detecting digital intrusions, its capabilities do not traverse into the physical realm to avert unauthorized access to server rooms.

q_ctrl_cattypes_physical_02_secp8

After a recent server outage, the company discovered that an employee accidentally unplugged the power cable from the server while grabbing some office supplies from the nearby shelf.

What security control did the company lack that led to the server outage?

Answers:

- Managerial
- Technical
- Operational
- ***Physical**

Explanation:

Physical controls such as alarms, gateways, locks, lighting, and security cameras deter and detect access to premises.

Managerial controls provide oversight of the information system. Examples could include risk identification or a tool allowing the evaluation and selection of other security controls.

Technical controls are the implementation of a system, such as hardware, software, or firmware. For example, firewalls, antivirus software, and OS access control models are technical controls. For the server's security, segregating that equipment from normal employee access is important.

Operational controls are implemented primarily by people. For example, security guards and training programs are operational controls.

q_ctrl_cattypes_preventive_01_secp8

Which of the following is an example of a preventative control type?

Answers:

- ***An advanced network appliance**
- Real-time monitoring alerts
- Intrusion detection systems
- Network monitoring applications

Explanation:

The easiest prevention control is an advanced network appliance, which is sometimes called an adaptive security appliance (ASA).

Examples of detective controls are intrusion detection systems (ISPs), network monitoring applications, collectors logs, and real-time monitoring alerts.

q_ctrl_cattypes_preventive_02_secp8

A company installed a new locking cabinet in the computer room to hold extra flash drives and other supplies.

Which type of security control did the company configure?

Answers:

- ***Preventive**
- Compensating
- Containment
- Deterrent

Explanation:

A preventive control physically or logically restricts unauthorized access. A system password and physical door lock are examples of preventive controls.

Rather than preventing an attack, a compensating control actively restores the functionality of systems via alternative methods, such as using a backup.

Containment does not refer to a security control type; it refers to a step in the incident management lifecycle for handling a threat.

A deterrent control may not physically or logically prevent access, but psychologically, it discourages an attacker from attempting an intrusion. For example, a warning sign is a deterrent control.

q_ctrl_cattypes_security_control_secp8

After an unauthorized access incident in the server room over the weekend, the IT department of a company decides to implement new security controls to deter similar future incidents.

Which of the following should they implement?

Answers:

- ***Placing visible signs indicating surveillance and severe penalties for unauthorized entry**
- Installing a network-connected smoke detector system in the server room
- Applying reflective window film to the server room windows
- Introducing a key control system for office desks

Explanation:

Placing visible signs indicating surveillance and severe penalties for unauthorized entry is an example of a physical control. This kind of deterrence can discourage unauthorized physical access by setting clear consequences.

Installing a network-connected smoke detector system in the server room is a physical control, but it's designed to detect fire hazards, not deter unauthorized access.

Applying reflective window film to the server room windows is a physical control aimed at enhancing privacy, but it doesn't deter unauthorized access once inside the building.

Introducing a key control system for office desks is a physical control, but it's more relevant for securing individual workstations and personal belongings than deterring unauthorized access to a server room.

q_ctrl_cattypes_technical_01_secp8

Given the need to prioritize cost-effective solutions for enhancing the company's cybersecurity posture, a global corporation's chief security officer (CSO) considers implementing technical controls over physical controls.

Which of the following options is a technical control?

Answers:

- ***Setting up a network intrusion detection system**
- Implementing a risk identification tool
- Conducting employee cybersecurity training
- Installing a building access control system

Explanation:

Setting up a network intrusion detection system is an example of a technical control. It involves hardware and software systems specifically designed to monitor and control the network's security, thus serving as a cost-effective solution for improving the cybersecurity posture.

A risk identification tool falls under the category of managerial controls. These controls oversee the information system and aid in selecting and implementing other security controls.

Employee cybersecurity training is an operational control. Operational controls are primarily human-centric and focus on procedures and responsibilities to maintain the security of the organization.

Installing a building access control system is a physical control. Physical controls manage access to premises and hardware, often incurring higher implementation and maintenance costs compared to technical controls.

q_ctrl_cattypes_technical_02_secp8

As part of enhancing its data protection strategy, a corporation's IT manager aims to ensure defense-in-depth by integrating a technical control alongside existing managerial and operational controls.

Which measure BEST exemplifies a technical security control according to the classification scheme?

Answers:

- ***Setting up a network intrusion detection system**
- Implementing a risk identification tool
- Conducting employee cybersecurity training
- Installing a building access control system

Explanation:

Intrusion detection systems represent a technical control involving hardware and software systems specifically designed to monitor and control the network's security. Network intrusion detection systems are naturally automated and technical, making this the best example of a technical control in the options given.

Implementing a risk identification tool falls under the category of managerial controls. These types of controls oversee the information system.

Conducting employee cybersecurity training is an operational control. Operational controls are primarily people-oriented.

Installing a building access control system is a physical control. Physical controls are measures that deter and detect access to premises and hardware.

q_ctrl_cattypes_technical_03_secp8

A company finds that employees are accessing streaming websites that are not being monitored for malware or viruses.

Which type of control can the network administrator implement to protect the system and keep the employees from viewing unapproved sites?

Answers:

- ***Technical**
- Operational
- Detective
- Corrective

Explanation:

A technical security control includes hardware or software mechanisms used to protect assets. Additionally, antivirus software, firewalls, and intrusion detection systems are examples of technical controls.

Operational security control characterizes a tangible item, preventing or detecting unauthorized access to physical spaces, systems, and assets.

A detective control identifies when incidents or vulnerabilities have occurred. For example, auditing and monitoring would be detective controls.

Corrective control actively responds to an incident, fixes it, and prevents it from happening again. Antivirus software exemplifies a corrective control.

1.3 Use the Simulator

As you study this section, answer the following questions:

- How do I complete simulation labs in this course?
- What features have been simulated in this environment?
- How will getting acquainted with the simulation environment help me acquire the necessary skills?

In this section, you will learn to:

- Read simulated component documentation and view components to make appropriate choices to meet the scenario
- Add and remove simulated computer components
- Change views and add simulated components
- Use the zoom feature to view additional image details
- Use the simulation interface to identify where simulated cables connect to the computer
- Attach simulated cables
- Configure a security appliance
- Install a security appliance

Key terms for this section include the following:

Term	Definition
Lab simulator	The lab simulator is a LabSim learning tool that presents a virtual environment that you can manipulate like an actual environment.
Lab tasks	The tasks necessary to complete the lab.
Navigation bar	A lab simulation feature used to change to a new location, such as a building, floor, or office.
Shelf	An area that contains hardware components that may be used in the simulation.
Exhibits	Additional information about the simulation environment that may be useful in completing the lab.

1.3.1 Use the Simulator (Demo Video)

Transcript:

TestOut's lab activities are key to your training. In this demonstration, I'm going to show you the components of the Lab Simulator so that you can successfully complete the activities in this course.

Some labs start out with an overview of the office. We'll talk about that a bit more later, but first, let's click on Hardware here in Office 2. The lab has four main areas. Over here, on the left, is the Scenario window. This window is very important. It describes the tasks that you're required to perform during a given lab activity. Typically, the items in this bullet list provide you with all the tasks that you'll be evaluated on. You'll be expected to perform these tasks correctly as you go through the lab. If you need more space while you're working, you can hide the Scenario window by clicking this button right here, and you can click it a second time to bring it back.
The main area where you'll do most of your work is called the Workspace. It includes all the items you'll work with and configure. For example, in this simulation, we have a computer, and we have wall plates with connectors for cable internet, the network, the telephone, and for AC power.
We also have the Shelf. This holds pieces of equipment organized by category. These are the objects you'll use to complete configuration tasks over in the Workspace. You can think of the Shelf as your inventory of spare parts or an online ordering system where you can order the parts that you need.
Now, let's go through the process of completing a lab. The first thing you need to do is read the Scenario. Read it very carefully because when you're done, you're evaluated on whether you did everything it asked you to.
You'll often need to examine objects within the Workspace more thoroughly. You can use this slider to zoom in and out. You can also use the zoom out and zoom in buttons or the drop-down list here.
Before we go any further, I need to point out that each object within the Workspace occupies a certain amount of space, which is denoted by the outline that's around each object. For example, this is the area for the computer, and this is the area for the wall plates. Within each of these windows, there are buttons that allow you to change the viewing perspective for that object.
I'm currently looking at the front of the computer. But let's that suppose I need to do some work on the back. I can come up here and click the Back button. If I need to see the front of the computer again—say, to power it on—I click the Front button. You'll notice that not all of the objects have multiple views. The computer does, but the wall plates only have one view—the Front view.
Let me sort of jump ahead for the sake of showing you something. I'm going to grab the mouse and keyboard from under Input Devices on the Shelf. You're not stuck with this layout in the Workspace. For example, I could grab the mouse and move it to the other side of the keyboard or move it back.

Once you've familiarized yourself with the items in the Workspace, you need to go over here—to the Shelf—and use the categories displayed to find the objects that are required to complete the Scenario.
If you're looking at an item on the Shelf, and you're not sure what it is, click on that object's Details link. For example, I can click the Details link here for the video cable. When I do, an overview of the cable is displayed. Notice that, for this object, we see the cable itself. We can see each connector, and I can look at the front, back, and top of each one. Also, notice that when I click an item on the Shelf, it's displayed down here in the Selected Component window, and we can access the same information using the Details links down here. Being able to view the details of a particular object on the Shelf is very useful because it helps you verify that the object you've selected is the correct one for the scenario's particular requirements.
And I should point out here that, with some objects, when you view their details, you see an additional tab called Specifications, which opens a window that provides even more information. Right now, we don't have an object on the Shelf or Workspace that has a Specification tab, but some do—for example, a motherboard will most likely have one. The information there is similar to what you might find in a user manual for that particular item.
Before we can work with an object in the Workspace, we have to first add it or connect it to an item that's already there. For example, to add this Monitor to the Workspace, we expand Monitors and then drag it over here. Notice that when I do, you see yellow lines appear, which tell me where I can drop the monitor in relation to the other objects that are already there. In this case, I want to add the monitor to the Workspace right next to the PC system itself, so I'm going to drop it right here. And now, it's there. Pretty simple!
Once the object is in the Workspace, I can manipulate it to accomplish the tasks in the Scenario. In this example, I need to look at the back of the monitor, and I need to use cables to connect this monitor to the computer system and the power outlet. Let's look at the back of the computer as well as the monitor to see all the connectors.
With this done, I now need to use the appropriate cables to connect these two devices together and connect the monitor to the power outlet. Let's go over here and expand Cables, and let's connect the monitor to the computer system using a video cable. I can click on the video cable so that it appears down here, in the Selected Component window. Let's grab each connector and add it to the appropriate port on the back of the computer and monitor. And let's quickly drag this connector to the monitor.
Notice here that there are three different ports available, and as I hover over each one, it's outlined in blue. That blue rectangle tells me that this is a potential place where I could connect this particular device. It doesn't mean it's the right connector, though—it just means that it's an option you could try. In this case, I'm dealing with an HDMI connector, so I need to make sure that I drag and drop it in an HDMI port—right here. I release the mouse, and now, one end of the connector is connected to the monitor. If we look at the connector's status down here, it's got one end connected to the monitor now. The other is still unconnected. And notice that when I did that, the Partial Connections window was displayed. You'll always see this window when one end of a cable is connected, but the other end isn't.
Now we need to connect the other end of the cable to the PC system. I'm going to click and drag. And just like with the other end of the cable, I need to pick the right port to connect it to. And there it is. Once again, the status is updated down here, in the Selected Component window. One end is connected to the computer and the other to the monitor.
Now, you might be wondering, "What happens if I drop this on an incompatible connector?" Well, when you do, an error is displayed down here, saying, "Hey, you can't connect that there." Let's go ahead and put it back on the correct port. And that's one way you can connect devices together using an item from the Shelf.
There's another way to do it as well, and that's to drag the cable directly from the Shelf and then drop it on the appropriate connector. In this case, I'm going to drag the power connector, and I'm going to drop it on the monitor's power socket. Notice that this cable has two different connectors—we have a female connector and a male connector. The simulator doesn't know which end of the cable I want to connect. It brings up a list of possibilities, and I have to tell it specifically which one I want to use. Let's go ahead and use the AC Power Female Connector. And now that the end is connected to the monitor itself, let's plug the other end into the AC wall plate.
Remember, I also moved our mouse and keyboard to the Bench from the Shelf. Let's plug those in really quick. First, I'll select the keyboard, find the USB on the back of the PC, and plug it in. Now let's do the same thing for the mouse.
Before we can use the computer or monitor, we obviously need to turn them on. Let's start with the computer. If you hover over the power button, you'll notice that it's highlighted in blue. I'm going to click it to turn it on. It's telling me that the monitor has no display—that's because it's not turned on. Let's go over to the monitor and power it on as well. Now I'll click on the Windows screen.
As you can see, we have a fully simulated Windows environment, though it does function in pretty much the same way a real Windows desktop would. For example, I can click on the Start button. When I do, all of the things that you'd expect to see in the Start menu are displayed. I can search for control panel and click on it. When I do, Control Panel pops up, just like it would on a real Windows system. And using the various links in Control Panel, I can go ahead and configure this simulated workstation. For example, I could go down here to Hardware and Sound, and I could use this link right

here to add a new printer to the system. The steps that you need to take within the simulation are the same ones you'd need to take on a real Windows system.
I do need to point out that as you go through the lab exercises, you'll see that not everything in the Windows interface is enabled. If a feature isn't necessary for a lab, it's not enabled, but all the components you need to complete the Scenario will be.
Now, while the computer is up and running, we can switch back to the Workspace and view the hardware by clicking up here.
In addition to moving objects from the Shelf to the Workspace, you can do just the opposite. You can take an object that's currently in the Workspace and return it back to the Shelf. Let's grab the keyboard and move it to the shelf. When I do, I get a message that says, "Hey, I can't be moved back to the Shelf because I'm still plugged in."
Before we can put this item back, we've got to unplug it. To unplug a cable or device, simply click on it, drag the end off, and drop it somewhere within the Workspace. For example, to unplug the cable for the keyboard, I click it, drag it off, and drop it—then its status changes to unconnected. At this point, the keyboard is unplugged, so I can drop it back over here on the Shelf. And please be aware that there are some items that you can't move to the Shelf, such as the wall plate connectors.
Now, once you've completed all the tasks in the Scenario, you're ready to submit the lab for evaluation. However, I recommend that before you do this, you go back over to your Scenario and quickly review all the tasks that you were expected to complete and double-check to make sure that everything was done as required. Once you've verified that you've done everything, go up here and click Score Lab. When you do, it's going to evaluate whether or not you did everything correctly.
Notice here that we have a list of tasks that I was required to perform for the Scenario—found here, under TASK SUMMARY. Over here, I have an icon that tells me whether or not I did each one. If I didn't do a particular task, I see a red X. If I did complete a task, it's denoted with a green checkmark. Notice that, in this scenario, I didn't do everything that was required. I initially had the keyboard on the Workspace, but I took it off, so it evaluated as incorrect. And because I unplugged the cable for the keyboard, I got it wrong.
Now, this last task is actually a multi-item task, meaning that I had to complete several different sub-tasks in order to get everything correct. To see what those sub-tasks were, I click on Show Details, and it tells me that I needed to turn on the monitor and the computer. I did both, so I got those correct.
And down here, under Task Summary, is the EXPLANATION. This part of the report provides step-by-step instructions for actually completing the tasks in the Scenario. There are, of course, many ways to complete some of these tasks.
Okay, now I'm going to pause the recording while I restart the lab.
Let's go back to Office 2 Hardware. If, at any time, you're in the middle of a lab and think to yourself, "I really messed this up," you can restart the lab without scoring it. This button over here will restart the lab from the beginning without giving you a score. Be aware, though, that if you do, you'll get this message telling you that any work you've done will be lost.
That's it for this demonstration. In this demo, we learned how to use the TestOut Lab Simulator.

1.3.2 Labsim Features (Demo Video)

Transcript:

In this demonstration, we're going to spend a few minutes talking about the lab exercises in this course. You may be familiar with LabSim already, but in this course, we're introducing a few new features, so you'll want to watch this demo.

When you start a lab exercise, you'll generally see a screen like this. We have a simulated a Kali Linux environment.
Let's take a look at it. If you want to complete the lab and get everything right, you should read through this scenario very carefully. It'll usually begin by setting up the situation. It's telling you what's happened and what we're trying to do in the demo. Down here, it'll give you the specific tasks that you need to complete. For this lab, we need to crack the password on the Support computer and then crack the password on the ZIP file.
As you go through this lab exercise, you need to make sure that you complete each and every one of these requirements. When you're done with the lab, each task is evaluated, and how many you completed determines your score.
So, as you're working through the lab exercise, it's a really good idea to mentally check and say, "Okay, did I do this one?" Yes. "Did I do this one?" Yes. If you've done everything listed here according to the scenario, you'll pass the lab exercise.
Now, our new lab features in this course are questions and copy and paste.

Let's start with questions. You may be required to perform a task and look for something specific. Over here, you'll see the Answer Questions tab. I have two questions that I'll need to answer while doing this lab. It's a very good idea to open this right after you read the scenario so you know what you're looking for while you perform the steps. This particular lab wants you to enter the password for the Linux Computer and the password for the protected.zip file.
Okay, now let's move on to the copy and paste feature. Let's open a terminal, and I'll type in **ifconfig.** Now let's say, for example, that I want to copy and paste this MAC address, here. I can just highlight it like this, come down here, and right-click. When I do, it'll paste the text right here. Be aware you don't get a menu that asks you to paste; it just pastes the text when you right-click.
If there's some text over in the scenario, let's say a long password, you can also copy that and paste it somewhere else. So, as soon as I highlight the text, it's copied to the clipboard. Come over here, right-click, and it's pasted in. Let's close the terminal.
Now, within each lab, not everything is actually enabled within the simulation, only the components that you'll need to complete the scenario. So you might click on something and find that it doesn't actually work, or it might not do what you think it should do.
But we're going to look at some of the things that do work in this lab. Over here, we have our Favorite programs. The top one, here, is the Terminal. I'll open that up again. Now, be aware that many of the programs that come with Kali Linux and are launched from the terminal. For example, if the lab says to use nmap to do a scan, I can do that from here. Let's check our IP address really quick. I can do that by typing **ifconfig** and pressing Enter.Right here, I can see my IP address of 192.168.0.45. We can also type **ip addr** andpress Enter to get our IP address.
Remember we launch nmap from the terminal; let's do that now. Just type in **nmap** and press Enter.When you do that, you get a list of parameters that you can use. Right here are some examples of how to use nmap. Now let's scan our subnet. For that, type **nmap 192.168.0.0/24** and press Enter.We just did a quick default nmap scan of our subnet, and we found all these devices that are live on the subnet.
Other programs are also launched from the terminal. To launch Metasploit, we'll type **msfconsole** and press Enter, and it's launched.
Okay, let's close the terminal and look at a few of the GUI tools that are simulated. We have Ettercap here. There will be labs that require you to use Ettercap in the course. Let's close that and open up the next one, Zenmap.
Zenmap is basically a GUI version of nmap. It's used to scan networks. Let's close this and go to the next one, Wireshark.
To start a scan, we select our interface and click the shark fin up here. Let's open the terminal. I'll do an **ifconfig** to confirm my IPaddress and ping this machine, so let's type in **ping 192.168.0.46** and press Enter.
Now let's go back to Wireshark and filter by ICMP, or our ping traffic, by typing in **icmp** in the filter field.You can see we're getting all this ICMP traffic in Wireshark.
Let's go ahead and close these.
Be aware that you may be required to open multiple programs or perform tasks on multiple systems. If this is the case, the scenario will tell you. This particular scenario only requires us to work on this one system, but there will be scenarios that will ask you to work on one, two, or more different systems. You might be asking, how do I change between systems? What you need to do is come up to the overview button, right here, and click on it. When you do, the entire office environment is displayed. The current system that you're working on is highlighted. Notice, down here, that the last system I was working on was named IT-Laptop. If you need to change to a different system somewhere else in the office, locate the appropriate system in the diagram, and then just click on it. For example, I'll click on the Office1 system. Now notice I'm in a different office, working on a completely different system. Let's go the IT-Laptop system.
Since this is a simulated environment, keep in mind that keyboard shortcuts might not work as expected. So, for example, if you want to enter Ctrl+Alt+Delete, you'll be sending that command to the actual system you're on, not the simulated system. Always look for text that lets you know how to perform the alternative to shortcuts; it's typically down here, at the bottom.
Now that we've looked at some of the new features, let's complete this lab. This lab tells us to crack some passwords with John the Ripper. I'll open the terminal and, at the prompt, I'll type in **cd /usr/share/john** and press Enter.Now let's see what's in the folder. I'll type in **ls** to do that.
One of the things I want to look at is the password.lst file. This will contain a list of passwords that John the Ripper will use to try to brute force and crack the system with. You can add to this list, delete passwords, and so on. To view the contents, I'll type in **cat password.lst** and press Enter, and you can see the list is displayed. This is a short list of passwords, so be aware that normally, this will contain tens of thousands of passwords. I'll type **cd** to get back to root, here,and then **clear** to clear the screen.
To crack the password on the system, I'll type in **john /etc/shadow** and press Enter.In this course, we'll learn that with Linux, the password is actually kept in the shadow file. Right up here, we can see that our password was cracked. As

you might remember, that was one of the questions we need to answer, so let's click on Answer Questions. I'll type in my answer and minimize that box. I'll clear this screen again.
Down here, it says once you've cracked the password, you can't crack it a second time, so let's see if that's true. I'll type in **john /etc/shadow** and pressEnter, and this time, it's telling me there are no password hashes left to crack.
It also tells me that the results are located in the john.pot file, so let's look in there to view the password, just in case we forgot to answer our question. For that, I'll type in **cat ./.john/john.pot,** press Enter, and here are the results. An alternative is to type in **john /etc/shadow/** --show', and you can see the password results again.
The next step says to crack the password of the protected.zip file on the IT-Laptop. Let's say, for some reason, we don't complete this part. I'll go back up to Answer Questions, type in an answer for the second password, and click on Score Lab, down here.
My results are displayed. I got 1 out of 4 complete, or 25%. I've been in the lab for about 30 minutes. Down here, it says I didn't correctly crack the password; however, I did correctly type in the right password as an answer. I didn't even attempt this next part, and I typed in some random answer right here. And down here is the real answer.
If I'm having issues, I can come down here, read this explanation, and review the correct steps I would need to take to complete the lab.
I can scroll down and click on the Done button; the grades are recorded in the gradebook.
That's it for this demonstration. In this demo, we talked about how to complete the lab exercises in this course. We looked at the scenario. We looked at the simulated Kali Linux environment. We discussed some new features in Labsim. We talked about how to complete tasks within the environment. We talked about how to switch between systems. We also talked about some key things you need to keep in mind when you're working in the simulated environment. We ended this demonstration by discussing how the labs are scored.

2.0 Threats, Vulnerabilities, and Mitigations

2.1 Understanding Attacks

As you study this section, answer the following questions:

- What motivates threat actors to attack?
- What protections can you implement against inside threat actors?
- Which three types of threat actors are most likely to have high levels of funding?
- Why are nation-state threat actors especially dangerous?
- What attack surfaces are inherent within a supply chain?

Key terms for this section include the following:

Term	Definition
Threat actor	A person or entity responsible for an event that has been identified as a security incident or as a risk.
Internal/external	The degree of access that a threat actor possesses before initiating an attack. An external threat actor has no standing privileges, while an internal actor has been granted some access permissions.
Level of sophistication/capability	A formal classification of the resources and expertise available to a threat actor.
Resources/funding	The ability of threat actors to draw upon funding to acquire personnel, tools, and development of novel attack types.
Service disruption	A type of attack that compromises the availability of an asset or business process.
Data exfiltration	The process by which an attacker copies data from a private network to an external network.
Disinformation	A type of attack that falsifies an information resource that is normally trusted by others.
Blackmail	Demanding payment to prevent the release of information.
Extortion	Demanding payment to prevent or halt some type of attack.

Fraud	Falsifying records, such as an internal fraud that involves tampering with accounts.
Hacker	Often used to refer to someone who breaks into computer systems or spreads viruses. Ethical hackers prefer to think of themselves as experts on and explorers of computer security systems.
Unauthorized hacker	A hacker operating with malicious intent.
Authorized hacker	A hacker engaged in authorized penetration testing or other security consultancy.
Unskilled attacker	An inexperienced attacker that typically uses tools or scripts created by others.
Hacktivist	A threat actor that is motivated by a social issue or political cause.
Advanced persistent threat (APT)	An attacker's ability to obtain, maintain, and diversify access to network systems using exploits and malware.
Nation-state actors	A type of threat actor that is supported by the resources of its host country's military and security services.
Organized crime	A type of threat actor that uses hacking and computer fraud for commercial gain.
Internal threat	A type of threat actor who is assigned privileges on the system that cause an intentional or unintentional incident.
Unintentional or inadvertent insider threat	A threat actor that causes a vulnerability or exposes an attack vector without malicious intent.
Shadow IT	Computer hardware, software, or services used on a private network without authorization from the system owner.
Vulnerable software	Weakness that could be triggered accidentally or exploited intentionally to cause a security breach.
Unsupported systems	Product life cycle phase where mainstream vendor support is no longer available.
Unsecure network	Configuration that exposes a large attack surface, such as through unnecessary open service ports, weak or no authentication, use of default credentials, or lack of secure communications/encryption.

Lure	An attack type that will entice a victim into using or opening a removable device, document, image, or program that conceals malware.
Supply chain	The end-to-end process of supplying, manufacturing, distributing, and finally releasing goods and services to a customer.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.1 Compare and contrast common threat actors and motivations.
	• Threat actors
	  ○ Nation-state
	  ○ Unskilled attacker
	  ○ Hacktivist
	  ○ Insider threat
	  ○ Organized crime
	  ○ Shadow IT
	• Attributes of actors
	  ○ Internal/external
	  ○ Resources/funding
	  ○ Level of sophistication/capability
	• Motivations
	  ○ Data exfiltration
	  ○ Espionage
	  ○ Service disruption
	  ○ Blackmail
	  ○ Financial gain
	  ○ Philosophical/political beliefs
	  ○ Ethical
	  ○ Revenge
	  ○ Disruption/chaos
	  ○ War
	
	2.2 Explain common threat vectors and attack surfaces.
	• Message-based
	  ○ Email
	  ○ Short Message Service (SMS)
	  ○ Instant messaging (IM)
	• Image-based
	  ○ File-based
	  ○ Removable device
	  ○ Vulnerable software
	  ▪ Client-based vs. agentless
	  ○ Unsupported systems and applications
	  ○ Unsecure networks
	  ▪ Wireless
	  ▪ Wired

	▪ Bluetooth ○ Open service ports ○ Default credentials ○ Supply chain ▪ Managed service providers (MSPs) ▪ Vendors ▪ Suppliers 2.3 Explain various types of vulnerabilities. • Supply chain ○ Service provider ○ Hardware provider ○ Software provider 2.4 Given a scenario, analyze indicators of malicious activity. • Application attacks ○ Privilege escalation 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. • Least privilege

2.1.1 Threat Actor Types (Lesson Video)

Transcript:

In cybersecurity, a threat actor is any individual or entity that carries out an attack. There are different types, attributes, and motivations of which to be aware. For example, a hacker trying to exploit a vulnerability has a different attack profile than

an organized crime group waging an assault on your network.
Before we look at the threat actors, it's important to understand their motivation. Let's look at the difference between opportunistic and targeted attacks. In an opportunistic attack, the threat actor usually tries to make easy money as fast as possible and with minimal effort. Because hiding their tracks and presence is time-consuming, they usually won't bother. They want to get in, grab the goods, and get out. Think of it as the smash-and-grab of cyberattacks.
Opportunistic attacks are unstructured and scan a wide range of systems for known vulnerabilities, such as old software, exposed ports, default configurations, etc. A common example of an opportunistic attack is ransomware.
A targeted attack, on the other hand, is much more dangerous and difficult to defeat. A targeted attack is structured and focused on a specific organization or person. The motivation for a targeted attack is data exfiltration. The data is often used for other purposes, such as espionage, blackmail, financial gain, or revenge.
Now, let's look at different threat actors. An insider threat is anyone who has received rights or privileges to access systems as part of their responsibilities to the organization. This can be employees, contractors, or even guests who may access the network temporarily. If the organization doesn't have good offboarding controls, there's the blurred case of former insiders, such as ex-employees now working at another company or dismissed and now harbor a grievance.
The main motivators for a malicious internal threat actor are revenge and financial gain. You must also assess the possibility that an insider threat may be working in collaboration with an external threat actor or group. A whistleblower is someone with an ethical motivation for releasing confidential information. While this could be classed as an internal threat in some respects, it's important to realize that whistleblowers can't themselves be threatened or labeled in any way that seems retaliatory or punitive.

Becoming an internal threat actor (insider) doesn't always require an employee to make a conscious decision to carry out an attack. Insider threats can also arise from unintentional sources. An unintentional or inadvertent insider threat is often caused by a lack of awareness or carelessness, such as users demonstrating poor password management. Without proper training, they'll continue to compromise the organization in ignorance.

Another example of an unintentional insider threat is the concept of shadow IT. Shadow IT refers to hardware or software purchased or introduced to the workplace without the sanction of the IT department and without going through a procurement and security analysis process.

Many employees want quick solutions and may download or use apps and services that aren't IT-approved. The problem of shadow IT is exacerbated by the proliferation of cloud services and mobile devices, which are easy for users to obtain. Shadow IT creates a new unmonitored attack surface for malicious adversaries to exploit.

Unless the IT department has approved the solution, they could pose a threat to security. The key takeaway is that insiders have easier access to company assets than an outsider trying to break in, making them a more dangerous threat.

External threat actors are commonly referred to as hackers. The term hacker describes an individual with the skills to gain access to computer systems through unauthorized or unapproved means. Originally, hacker was a neutral term for a user who excelled at computer programming and computer system administration. However, the term gradually became associated with illegal or malicious system intrusions. The terms unauthorized (previously known as black hat) and authorized (previously known as white hat) are used to distinguish these motivations. It's also helpful to define additional threat actors based on different motivations and capabilities.

An unskilled attacker is someone who uses hacker tools without necessarily understanding how they work or having the ability to craft new attacks. Unskilled attacks might have no specific target or any reasonable goal other than gaining attention or proving technical abilities. Today, many threat actors are now likely to work as part of a team or group. The collaborative team effort means that these threat actors are able to develop sophisticated tools and novel strategies.

A hacktivist group, such as Anonymous, WikiLeaks, or LulzSec, uses cyber weapons to promote a political agenda. Hacktivists might attempt to use data exfiltration to obtain and release confidential information to the public domain, perform service disruption attacks, or deface websites to spread disinformation. Political, media, and financial groups and companies are most at risk of becoming a target for hacktivists, but environmental and animal advocacy groups may target companies in a wide range of industries.

In many countries, cybercrime has overtaken physical crime in terms of the number of incidents and losses. Organized crime can operate across the internet from a different jurisdiction than its victim, increasing the complexity of prosecution. Criminals will seek any opportunity for profit, but typical activities are financial fraud—against individuals and companies—and blackmail/extortion.

Most espionage is thought to be pursued by state actors, but it's not inconceivable that a business competitor might use cyber espionage against its competitors. Such attacks could aim at theft, disrupt a competitor's business, or damage their reputation. Competitor attacks might be facilitated by employees who have recently changed companies and bring insider knowledge with them.

Most nation-states have developed cybersecurity expertise and will use cyber weapons to achieve military and commercial goals. Nation-state actors have been implicated in many attacks, particularly on energy, health, and electoral systems. The goals of state actors are primarily disinformation and espionage for political or military strategic advantages, but it's known for countries to target companies for financial gain.

The term advanced persistent threat (APT) was coined to understand the behavior underpinning modern types of cyber adversaries. Rather than think in terms of systems being infected with a virus or Trojan, an APT refers to the ability of an adversary to achieve an ongoing compromise of network security—to obtain and maintain access—using various tools and techniques.

Well, that's it for this lesson. In this lesson, we discussed different types of attacks, such as opportunistic and targeted attacks. We also talked about different threat actors, including insiders, hackers, organized crime, competitors, and nation-states.

2.1.2 Threat Actor Types Facts

Historically, cybersecurity techniques relied on the identification of static known threats, such as viruses or rootkits, Trojans, botnets, and exploits for specific software vulnerabilities. It is relatively straightforward to identify and scan for these types of threats with automated software. Unfortunately, adversaries were able to develop means of circumventing this type of signature-based scanning.

The sophisticated nature of modern cybersecurity threats requires the creation of profiles of threat actor types and behaviors. This analysis involves identifying the attributes of threat actors' location, capability, resources/funding, and motivation. To evaluate these factors, you must be able to evaluate the sources of threats or threat actors. This topic will help you to classify and evaluate the motivation and capabilities of threat actor types so that you can assess and mitigate risks more effectively.

This lesson covers the following topics:

- Attributes of threat actors
- Motivations of threat actors
- Types of threat actors

Attributes of Threat Actors

Attribute	Description
Internal/external	Internal/external refers to the degree of access a threat actor possesses before initiating an attack. An external threat actor has no account or authorized access to the target system. A malicious external threat must infiltrate the security system using unauthorized access, such as breaking into a building or hacking into a network. Note that an external actor may perpetrate an attack remotely or on-premises. It is the threat actor that is external rather than the attack method. Conversely, an internal/insider threat actor has been granted permissions on the system. This typically means an employee, but insider threats can also arise from contractors and business partners.
Sophistication/capability	The level of sophistication/capability refers to a threat actor's ability to use advanced exploit techniques and tools. The least capable threat actor relies on commodity attack tools that are widely available. More capable actors can fashion new exploits in operating systems, applications software, and embedded control systems. At the highest level, a threat actor might use non-cyber tools such as political or military assets.
Resources/funding	A high level of capability must be supported by resources/funding . Sophisticated threat actor groups need to be able to acquire resources, such as customized attack tools and skilled strategists, designers, coders, hackers, and social engineers. The most capable threat actor groups receive funding from nation-states and organized crime.

Motivations of Threat Actors

Motivation is the threat actor's reason for perpetrating the attack. A malicious threat actor could be motivated by greed, curiosity, or some grievance, for instance. Threats can be characterized as structured/targeted or unstructured/opportunistic, depending on how widely an attack is perpetrated. For example, a criminal gang attempting to steal customers' financial data from a company's database system is a structured, targeted threat. An unskilled hacker launching some variant of the "I Love You" email worm sent to a stolen mailing list is an unstructured, opportunistic threat.

A threat actor with malicious motivation can be contrasted with an accidental or unintentional threat actor. An unintentional threat actor represents accidents, oversights, and other mistakes. To help analyze motivations, it is first useful to consider the general strategies that a threat actor could use to achieve an objective:

- **Service disruption** — prevents an organization from working as it does normally. This could involve an attack on their website or using malware to block access to servers and employee workstations. Service disruption can be an end in itself if the threat actor's motivation is to sow chaos or gain revenge. Service

disruption can be used as a blackmail threat, or it can be used as a tactic in the pursuit of some different strategic objective.

- **Data exfiltration** — transfers a copy of some type of valuable information from a computer or network without authorization. A threat actor might perform this type of theft because they want the data asset for themselves, so they can exploit its loss as blackmail or sell it to a third party.
- **Disinformation** — falsifies some type of trusted resource, such as changing the content of a website, manipulating search engines to inject fake sites, or using bots to post false information to social media sites.

You can relate these strategies to how they affect the CIA triad: data exfiltration compromises confidentiality, disinformation attacks integrity, and service disruption targets availability.

Chaotic Motivations

In the early days of the internet, many service disruptions and disinformation attacks were perpetrated with the simple goal of causing chaos. Hackers might deface websites or release worms that brought corporate networks to a standstill for no other reason than to gain credit for the hack.

This type of vandalism, for its own sake, is less prevalent now. Attackers might use service disruption and disinformation to further political ends, or nation-states might use it to further war aims. Another risk is threat actors motivated by revenge. Revenge attacks might be perpetrated by an employee, former employee, or any external party with a grievance.

Financial Motivations

As hacking and malware became more sophisticated and better commodified, the opportunities to use them for financial gain grew quickly. If an attacker is able to steal data, they might be able to sell it to other parties. Alternatively, they might use an attack to threaten the victim with blackmail or extortion or to perpetrate fraud:

- Blackmail is demanding payment to prevent the release of information. A threat actor might have stolen information or created false data that makes it appear as though the target has committed a crime.
- Extortion is demanding payment to prevent or halt some type of attack. For example, a threat actor might have used malware to block access to an organization's computers and demand payment to unlock them.
- Fraud is falsifying records. Internal fraud might involve tampering with accounts to embezzle funds or inventing customer details to launder money. Criminals might use disinformation to commit fraud, such as posting fake news to affect the share price of a company, promoting pyramid schemes, or creating fake companies.

Political Motivations

A political motivation means that the threat actor uses an attack to bring about some type of change in society or governance. This can cover a very wide range of motivations:

- An employee acting as a whistleblower because of some ethical concern about the organization's behavior.
- A campaign group disrupting the services of an organization that they believe acts in contradiction to their ethical or philosophical beliefs.
- A nation-state using service disruption, data exfiltration, or disinformation against government organizations or companies in another state in pursuit of war aims.

Nation-states commonly perpetrate espionage and disinformation attacks against one another, whether or not they are at war. In cybersecurity, espionage is a type of data exfiltration aimed at learning secrets rather than selling them or using the theft for blackmail.

There is also the threat of commercial espionage, where a company attempts to steal the secrets of a competitor.

Types of Threat Actors

It can also be helpful to evaluate the risk that well-known threat actor types or profiles pose to a business, given awareness of the general strategies and motivations. Note that the detailed process of analyzing the threat posed by a particular actor or adversary group is described as threat modeling.

Hackers

Hacker describes an individual with the skills to gain access to computer systems through unauthorized or unapproved means. Originally, *hacker* was a neutral term for a user who excelled at computer programming and computer system administration. Hacking into a system was a sign of technical skill and creativity that gradually became associated with illegal or malicious system intrusions. The terms unauthorized (previously known as black hat) and authorized (previously known as white hat) are used to distinguish these motivations. An authorized hacker always seeks authorization to perform penetration testing of private and proprietary systems.

Unskilled Attackers

An unskilled attacker is someone who uses hacker tools without necessarily understanding how they work or having the ability to craft new attacks. Unskilled attacks might have no specific target or any reasonable goal other than gaining attention or proving technical abilities.

Hacker Teams and Hacktivists

The historical image of a hacker is that of a loner, acting as an individual with few resources or funding. While the "lone hacker" remains a threat that must be accounted for, threat actors are now likely to work as part of a team or group. The collaborative team effort means that these threat actors are able to develop sophisticated tools and novel strategies.

A *hacktivist group* , such as Anonymous, WikiLeaks, or LulzSec, uses cyber weapons to promote a political agenda. Hacktivists might attempt to use data exfiltration to obtain and release confidential information to the public domain, perform service disruption attacks, or deface websites to spread disinformation. Political, media, and financial groups and companies are most at risk of becoming a target for hacktivists. Still, environmental and animal advocacy groups may target companies in a wide range of industries.

Nation-State Actors

Most nation-states have developed cybersecurity expertise and will use cyber weapons to achieve military and commercial goals. The security company Mandiant's APT1 report into Chinese cyber espionage units, shaped the language and understanding of cyber-attack lifecycles.

The term advanced persistent threat (APT) was coined to understand the behavior underpinning modern types of cyber adversaries. Rather than think in terms of systems being infected with a virus or Trojan, an APT refers to the ability of an adversary to achieve an ongoing compromise of network security—to obtain and maintain access—using various tools and techniques.

Nation-state actors have been implicated in many attacks, particularly on energy, health, and electoral systems. The goals of state actors are primarily disinformation and espionage for strategic advantage. Still, it is known for countries—North Korea is a good example—to target companies for financial gain.

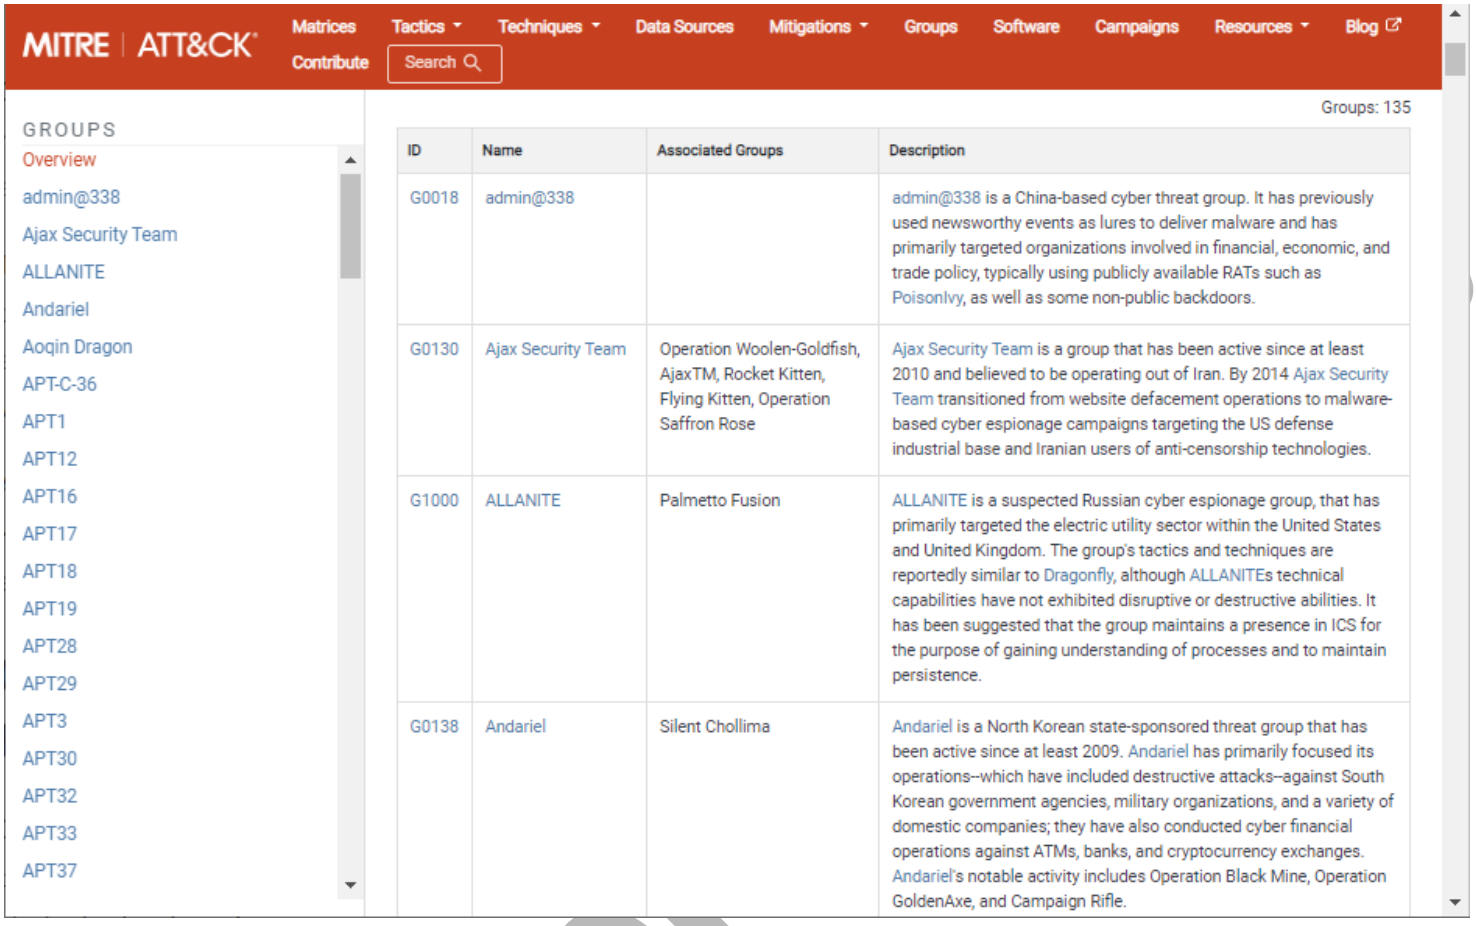

Researchers such as The MITRE Corporation report on the activities of organized crime and nation-state actors. (Screenshot © 2023 The MITRE Corporation. This work is reproduced and distributed with the permission of The MITRE Corporation.)

State actors will work at arm's length from the national government, military, or security service that sponsors and protects them, maintaining "plausible deniability." They are likely to pose as independent groups or even as hacktivists. They may wage false flag disinformation campaigns that try to implicate other states.

Organized Crime and Competitors

In many countries, cybercrime has overtaken physical crime in terms of the number of incidents and losses. Organized crime can operate across the internet from a different jurisdiction than its victim, increasing the complexity of prosecution. Criminals will seek any opportunity for profit, but typical activities are financial fraud—against individuals and companies—and blackmail/extortion.

Most espionage is thought to be pursued by state actors, but it is not inconceivable that a rogue business might use cyber espionage against its competitors. Such attacks could aim at theft, disrupt a competitor's business, or damage their reputation. Competitor attacks might be facilitated by employees who have recently changed companies and bring insider knowledge with them.

Internal Threat Actors

Many threat actors operate externally from the networks they target. An external actor has to break into the system without having any legitimate permissions. An internal threat (or insider threat) arises from an actor identified by the organization and

granted some type of access. Within this group of internal threats, you can distinguish insiders with permanent privileges, such as employees, from insiders with temporary privileges, such as contractors and guests.

There is the blurred case of former insiders, such as ex-employees now working at another company or who have been dismissed and now harbor a grievance. These can be classified as internal threats or treated as external threats with insider knowledge and possibly some residual permissions if effective offboarding controls are not in place.

The main motivators for a malicious internal threat actor are revenge and financial gain. Like external threats, insider threats can be opportunistic or targeted. An employee who plans and executes a campaign to modify invoices and divert funds is launching a structured attack. An employee who tries to guess the password on the salary database a couple of times, having noticed that the file is available on the network, is perpetrating an opportunistic attack. You must also assess the possibility that an insider threat may be working in collaboration with an external threat actor or group.

A whistleblower is someone with an ethical motivation for releasing confidential information. While this could be classed as an internal threat in some respects, it is important to realize that whistleblowers making protected disclosures—such as reporting financial fraud through an authorized channel—cannot themselves be threatened or labeled in any way that seems retaliatory or punitive.

Insider threats can also arise from unintentional sources. Unintentional or inadvertent insider threat is often caused by a lack of awareness or carelessness, such as users demonstrating poor password management. Another example of unintentional insider threat is the concept of shadow IT , where users purchase or introduce computer hardware or software to the workplace without the sanction of the IT department and without going through a procurement and security analysis process. The problem of shadow IT is exacerbated by the proliferation of cloud services and mobile devices, which are easy for users to obtain. Shadow IT creates a new unmonitored attack surface for malicious adversaries to exploit.

2.1.3 General Attack Strategy (Lesson Video)

Transcript:

Let's talk about the general attack strategy threat agents use to conduct an exploit.

Although there are a multitude of ways to carry out an exploit, many threat agents will actually follow the same general attack strategy. Understanding the methodology behind common attacks can help you defend your information assets better.
The first step in most attacks is some type of reconnaissance. The attacker needs to probe the target computer system or target network for vulnerabilities. They're also trying to gather information. They want to know about system hardware and network configurations. They want to know about individual users and their susceptibility. They'll use different tools to find different information.
An agent conducting a social engineering attack needs to gather information about the users on the system. They might dumpster dive, sifting through garbage to try to find critical information, such as passwords that have been written down. Any information about network users, systems, and processes can be useful to an attacker.
Another tactic is to call individual users and intimidate them enough to give out important information. An attacker will try to sound like someone who has authority in the organization. They'll threaten employees with discipline to acquire sensitive information.
An alternative to the intimidation approach is the sympathetic approach. The attacker will pretend to be a fellow employee at the organization who is in a desperate situation that can only be rectified by certain information.
There are variations on these methodologies, of course. For example, it is possible for someone to use aggressive emails, texts, or messages on social media to gain information.
If an agent is conducting a technical exploit, then the tools used will be a little bit different. A very common tactic is to use the port scan with the utility. The nmap utility is a security scanner that hackers use a lot. It scans a system to find out what ports are open. An open port on a system represents a potential hole that can be exploited.
A ping sweep is another favorite attack tool. The attacking system sends a ping, or an ICMP echo request, to a range of IP addresses say 192.168.1.1 to 192.168.1.254. Any valid IP address configured to respond to ping requests will

respond with an ICMP echo reply packet. This tells an attacker which IP addresses are active and available on the network.
If the reconnaissance goes well and the attacker has a lot of information to work with, the next step is breaching the system.
The way the breach is conducted depends on the type of exploit involved. If it's a social engineering exploit, they're going to use the data gathered from the user, such as their username and password. Alternatively, they could get a user to do something for them.
If the exploit is a technical exploit, then different processes will be used. The attacker might try to cause an exception error with a buffer overflow that gives them access to the system. They might identify an open port on a network server and try to exploit it. They might try to use password-cracking software to crack a password.
Once the system has been breached, the next step is to escalate privileges. Most breaches provide the attacker with a limited amount of access to the system, since most user accounts don't have high-level privileges. Even if they used a technical exploit, an attacker will probably only have a low level of access. That's not good enough. They want escalated privileges to the system that they're attacking. There are a variety of ways to do this, but many involve exploiting weaknesses in the system's design.
In addition to escalating privileges, attackers will also want to create an easy way to access the system a second time. This is called a backdoor. It makes it so they won't have to go through the reconnaissance and breaching process the next time they want to get in. There are all kinds of ways to set up a backdoor.
For example, an attacker could set up their own user account. They may even create a user account that looks the same as a user account that already exists. Through social engineering, they may have identified the user name of a particular person in the organization. And with some operating systems, you can create a username and then put a space at the end of it. The operating system will recognize it as a completely different user account, but most administrators will see it as identical to the legitimate user account.
The next phase in the attack is the stage. This step is optional; if the exploit was designed to steal information like credit card numbers, then an attacker is not going to be concerned with spreading the attack to other systems. If, however, the attack is an internet worm that's going to be spread all over the world, then the attacker wants to use the system that they've breached to spread the attack. In other words, the breached system becomes a staging ground for mounting attacks on additional systems.
The last step is to actually exploit. The attacker tries to gain whatever they initially attacked the system for. It could be customer information, ID numbers, credit card numbers, or embarrassing personal information. The goal of the exploit may not even be to steal; it might seek to create so much traffic on a particular system that the system can no longer function normally and provide services. In this way, the attack is actually denying services, rather than stealing anything. An exploit could also seek to corrupt or modify information. Of course, the actual exploit conducted depends on the motivation of the threat agent. An attacker could be trying to steal information for financial gain or trying to get back at an organization for a perceived wrong. They might be a cyber terrorist trying to bring down a critical infrastructure system.
Understanding general attack strategies will help you, as a security administrator, to better identify when attacks are occurring and defend your system against them. It can also help you identify weak spots in your system so you can prevent attacks altogether.
To quickly review, the steps in a general attack strategy are reconnaissance, breaching, escalating privileges, staging, and exploiting.

2.1.4 General Defense Strategy (Lesson Video)

Transcript:

As a security administrator, you're responsible for protecting your information from an almost limitless number of attacks and exploits. To guard your information from attacks, you need to formulate a general defense strategy using several fundamental security principles.

The first principle we need to look at is layering. Your defenses need to be layered. This means using multiple strategies to protect information.
For example, in order to log in to a computer system, by default, all you have to provide is the user name and password. Well, that's good--but what if the user name and password get compromised? The attacker who obtained the user name and password now has access to whatever privileges that user account has within the system.

But what if you required a fingerprint that's read via a biometric scanner in addition to a user name and password? With this extra security precaution, we create a layered defense for that particular system. Someone could come along and compromise a particular user's user name and password on that system, but they still can't log in because they don't have that user's fingerprint.

The one thing you need to remember when you're dealing with most attackers is that they're looking for low-hanging fruit; in other words, they're lazy. Let's say you have a choice of three different systems to try to compromise, and they're all worth about the same. Two of them are extremely difficult, with layered defenses that will take a lot of work and risk to compromise. The third system is really lax on their security.

Which one are you going to attack? The third one! So, by using layered defenses, you make sure that your system or network is not a target of opportunity. Is that going to deter a determined attacker, like a disgruntled employee or a cyber terrorist? No. To them, you're still a target of opportunity because they're not focused on anybody else--they are focused on you. But using layered defenses still makes it more difficult to access your system.

The next part of your overall defensive strategy is minimizing user access. You need to protect your information from the inside and the outside. Basically, each user should have access to the information they need to do their job and no more. This is called the principle of least privilege. One problem with the principle of least privilege is that some operating systems give users access to only what they have been specifically granted privileges to, while other operating systems basically give everybody access to everything unless they are specifically denied access. What you need to do is look at the systems where the information you're protecting is stored and determine whether the operating system is the first type or the second type.

Here's an example: We have three users who use this particular server. It's a Linux server. We have Fred, Mary, and Bob. Fred needs access to his home directory on the server, which is located in the /Home/Fred directory. Therefore, according to the principle of least privilege, Fred should have access to only this directory on that server, and no others. Mary needs to have access to her home directory right here, /Home/Mary. She is also working on a project that's developing a new widget for the company. She needs access a special directory named Widgets, too. Additionally, she's been granted access to a shared area right here on the server hard drive, where multiple users can share their information. So she is given that level of access. Bob, over here, has access to his home directory, /Home/Bob, and he has also been granted access to the shared directory /Shared.

This is a very common setup in the real world, but there's a problem with it: Individual users have home directories that are protected and inaccessible to anybody else. But then the administrator sets up shared directories, and it becomes much easier for the principle of least privilege to be violated. The problem isn't with the directory itself; the problem is with what people do with that shared directory. From a productivity standpoint, shared directories work well because they allow multiple people to collaborate on documents and work together efficiently. A user can put a file into the shared folder, and anyone who has access to the shared folder can open the file and do whatever work they need to.

A problem that can occur with shared directories, however, is users adding sensitive files to the shared folder. It can happen that users that have access to the shared folder should not have access to certain files that have been added to the shared folder. Shared folders ruin the strict privilege denial system.

The next component that you need to include in your overall defensive posture is randomness. The problem here is that human beings tend to be creatures of habit. We establish predictable patterns of behavior. For example, you probably take the same route to work at the same time each day. Once you're at work you probably follow the same daily routine checking emails, attending meetings, writing reports, completing tasks and the same weekly routine. This is bad from an information security standpoint because attackers can identify these personal patterns and then use them to their advantage.

As the security administrator for a system, you need to randomize your personal habits as part of your defensive strategy. If someone wants to break into your system and they're using a social exploit combined with a technical exploit, they might observe what time you come to work, what time you go on break, what time you run backups, etc. Whatever information they can glean during the reconnaissance phase of an attack helps them determine where vulnerabilities lie.

Another good thing to do is to randomize password change intervals. Security administrators usually make it so that every user has to change their password after 90 days. Well, what does that tell an attacker? They know that if they manage to compromise a user name and password, they've got a certain amount of days while that password is still valid. But if you set up your password change intervals to be 90 days once, then 30 days the next time, then 60 days, then 15 days, etc., then the attacker does not know how long that password is good for.

You should also try to obscure information. In other words, you want to make the reconnaissance process as unreliable and unpredictable as possible. You don't want an attacker to deduce behavioral trends and gather information about your system. One common recon tactic is to use ping sweeps. One thing you can do to obscure information is to make it so ping sweeps don't work. You do this by configuring a host-based firewall on every system and configuring that firewall not to respond to ICMP Echo requests. Basically, you make the systems un-pingable. If you disable ICMP Echo requests

in the host-based firewall, they won't respond. You obscure how many systems are in your network, as well as which IP addresses are assigned.

The last principle of a defensive posture is to keep things simple. Many administrators want to implement all the latest and most complex security measures to protect their information. That's good, but it becomes a problem if defenses are so complex that they become unusable. It is better to keep a simple defensive mechanism. Here, simple means easy to understand. If you don't fully understand how to manage a type of defense, don't use it until you are well-trained in how it works. You can't make sure a security system is doing its job if you don't know what it's supposed to do. You need to know how to interpret the output from the system and how to fix problems when they occur.

As a security administrator, you need to implement the concepts that we talked about in your defensive strategy in order to create a system that's much more difficult to breach or exploit.

In this lesson, we covered the various components that comprise an overall defensive strategy. We talked about layered defenses, the principle of least privilege, varied defense mechanisms, randomizing behaviors, and keeping things simple.

2.1.5 Attack and Defense Strategy Overview

Understanding the methodology behind common attacks can help you better defend your assets.

This lesson covers the following topics:

- Attack strategies
- Defense methodologies

Attack Strategies

General attack strategies incorporate some or all of the techniques explained in the following table.

Strategy	Description
Perform reconnaissance	*Reconnaissance* is the process of gathering information about an organization, including: • System hardware information • Network configuration • Individual user information
Use social engineering	*Social engineering* is the process of manipulating others into providing sensitive information. Social engineering tactics include: • Intimidation • Sympathy
Use technical approaches	A *technical* approach to obtaining information includes using software or utilities to find vulnerabilities in a system. Methods often used by hackers are: • Port scan • Ping sweep
Breach the system	A *breach* is the penetration of system defenses. It is often achieved by using information gathered through reconnaissance.

Strategy	Description
Escalate privileges	*Escalating privileges* is a primary objective of an attacker. Once an attacker has breached the system, obtaining higher privileges allows the attacker to access more information and gain greater control within the system.
Create a backdoor	*Creating a backdoor* is an alternative method of accessing an application or operating system for troubleshooting. Hackers often create backdoors to exploit a system without being detected.
Stage computers	*Staging* a computer involves preparing it to perform additional tasks in the attack, such as installing software designed to attack other systems. This is an optional step.
Exploit vulnerabilities	An *exploitation* takes advantage of known vulnerabilities in software and systems. Once a vulnerability has been exploited, an attacker can often: • Steal information • Deny services • Crash systems • Modify/alter information

Defense methodologies

General defense methodologies include the following items:

Methodology	Description
Layering	*Layering* involves implementing multiple security strategies to protect the same asset. *Defense in depth* or *security in depth* is based on the premise that no single layer is completely effective in securing assets. The most secure system/network has many layers of security and eliminates single points of failure.
Principle of least privilege	The *principle of least privilege* states that users or groups are given only the access they need to do their jobs and nothing more. When assigning privileges, be aware that it is often easier to give a user more access when it is needed than to take away privileges that have already been granted.
Variety	Defensive layers should incorporate a variety of methods. Implementing multiple layers of the same defense does not provide adequate protection against attacks.
Randomness	*Randomness* in security is the constant change in personal habits and passwords to prevent predictable behavior.
Simplicity	Security measures should provide protection but not be so complex that it is difficult to understand and use.

2.1.6 Attack Surfaces (Lesson Video)

Transcript:

The attack surface is all the points at which a malicious threat actor could try to exploit a vulnerability. To evaluate the attack surface, you need to consider the attributes of threat actors that pose the most risk to your organization.

From a threat actor's perspective, each part of the attack surface represents a potential vector for attempting an intrusion. A threat vector is the path that an attacker uses to execute a data exfiltration, service disruption, or disinformation attack. Sophisticated threat actors will make use of multiple vectors. In this lesson, we'll look at several of those vectors.

Vulnerable software contains a flaw in its code or design that can be exploited to circumvent access control or to crash the process. Typically, vulnerabilities can only be exploited in specific circumstances and are often patched by the vendor. However, because of the complexity of modern software and the speed with which new versions must be released to market, almost no software is free from vulnerabilities. Organizations that might not have an effective patch management system are exposed to this commonly exploited threat vector.

Vulnerable software continues to be an effective threat vector due to the continued use of unsupported systems and applications. An unsupported system is one where its vendor no longer develops updates and patches. Unless the organization is able to patch the faulty code itself, these services and apps will be highly vulnerable to exploits. One of the best ways to mitigate this risk is to isolate these systems.

Scanning software helps organizations automate the discovery and classification of software vulnerabilities. This scanning software can be implemented as a client-based agent, where the agent runs as a scanning process installed on each host and reports to a management server. Alternatively, the vulnerability management product might use agentless techniques to scan a host without requiring any sort of installation.

To execute malicious code on a system, a threat actor must be able to run exploit code on the system or over a network. Let's look at some specific threat vectors that can expose unsecured networks.

Direct Access — when the threat actor uses physical access to the site or system to perpetrate an attack.

Wired Network — when a threat actor with access to the site attaches an unauthorized device to a physical network port, and the device is permitted to communicate with other hosts.

Remote and Wireless Network—when the attacker obtains credentials for a remote access or wireless connection to the network or cracks the security protocols used for authentication.

Cloud Access—many companies now run part or all of their network services via internet-accessible clouds. The attacker only needs to find one account, service, or host with weak credentials to gain access.

Some additional network vectors include the following:

Bluetooth Network—a threat actor can exploit a vulnerability or misconfiguration to transmit a malicious file to a user's device over the Bluetooth personal area wireless networking protocol.

Default Credentials—the attacker gains control of a network device or app because it has been left configured with a default password. Default credentials are likely to be published in the product's setup documentation or are otherwise easy to discover.

Open Service Port—the threat actor is able to establish an unauthenticated connection to a logical TCP or UDP network port. The server will run an application to process network traffic arriving over the port. The software might be vulnerable to exploit code or to service disruption.

A lure is something superficially attractive or interesting that causes its target to want it, even though it may be concealing something dangerous, like a hook. In cybersecurity terms, when the target opens the file bait, it delivers a malicious payload hook that'll typically give the threat actor control over the system or perform service disruption.

A common lure attack involves removable devices. An attacker conceals malware on a USB thumb drive or memory card and tries to trick employees into connecting the media to a device. In a drop attack, the threat actor simply leaves infected USB sticks near offices with the expectation that at least one employee will pick one up and plug it into a computer. They can contain executable files concealing malware, word processing documents or PDF files with embedded malicious code, or image files with exploit code concealed in the image.

File-based lures need a mechanism to deliver the file and a message that'll trick a user into opening the file on their computer. Consequently, any features that allow direct messaging to users must be considered as part of the potential attack surface. Message-based threats involve malicious activities carried out through different communication channels, such as email, short message services (SMS), and instant messaging.

A common vector involves sending email messages with malicious file attachments or fake URLs, commonly referred to as phishing. The attacker also uses social engineering techniques to persuade or trick the user into opening the attachment or clicking on the URL.

Another vector utilizes the Short Message Service (SMS) to send a file or a link to a mobile device using the text messaging handler built into smartphone firmware. This practice is called smishing, and it's difficult to monitor because it's operated by the handset or subscriber identity module (SIM) card provider.

Attackers can also exploit instant messaging (IM) platforms to deliver malicious links, files, or messages, attempting to compromise devices or steal sensitive data.
Web and Social Media sites provide yet another vector. Malware may be concealed in files attached to posts or presented as downloads. An attacker may compromise a site so that it automatically infects vulnerable browser software (a drive-by download).
A supply chain is an end-to-end process of designing, manufacturing, and distributing goods and services to customers. Rather than attack a target organization directly, a threat actor may seek ways to infiltrate it via companies in its supply chain.
Ensuring reliable sources of equipment and software is called procurement management. In procurement management, it's helpful to distinguish several types of relationships, such as suppliers, vendors, and business partners.
Suppliers obtain products directly from a manufacturer to sell to other businesses. Vendors receive products from suppliers to sell to retail businesses or directly to customers. The term business partners implies a closer relationship where two companies share closely aligned goals and marketing opportunities. Each of these companies has its own supply chain.
The supply chain breadth and complexity expose organizations to a huge attack surface. For example, a computer motherboard supply chain can include a chip manufacturer, firmware code developer, OEM reseller, courier delivery company, and administrative staff responsible for provisioning the computing device to the end user. Anyone with the time and resources to modify the computer's firmware could create backdoor access. The same applies to any computer or network hardware, software, or service.
Well, that's it for this lesson. In this lesson, we introduced you to multiple attack surface vectors including Vulnerable software vectors, Network Vectors, Lure-based Vectors, Message-based Vectors, and Supply Chain Vectors. Understanding the methods by which threat actors infiltrate networks and systems is essential for you to assess the attack surface of your networks and deploy controls to block attack vectors.

2.1.7 Attack Surfaces Facts

The attack surface is all the points at which a malicious threat actor could try to exploit a vulnerability. Any location or method where a threat actor can interact with a network port, app, computer, or user is part of a potential attack surface. Minimizing the attack surface means restricting access so that only a few known endpoints, protocols/ports, and services/methods are permitted. Each of these must be assessed for vulnerabilities and monitored for intrusions.

This lesson covers the following topics:

- Attack surface and threat vectors
- Vulnerable software vectors
- Network vectors
- Lure-based vectors
- Message-based vectors
- Supply chain attack surface

Attack Surface and Threat Vectors

An organization has an overall attack surface. You can also assess attack surfaces at more limited scopes, such as that of a single server or computer, a web application, or employee identities and accounts. To evaluate the attack surface, you need to consider the attributes of threat actors that pose the most risk to your organization. For example, the attack surface for an external actor should be far smaller than that for an insider threat.

From a threat actor's perspective, each part of the attack surface represents a potential vector for attempting an intrusion. A threat vector is the path that a threat actor uses to execute a data exfiltration, service disruption, or disinformation attack. Sophisticated threat actors will make use of multiple vectors. They are likely to plan a multistage campaign rather than a single "smash and grab" type of raid. Highly capable threat actors will be able to develop novel vectors. This means the threat actor's knowledge of your organization's attack surface may be better than your own. Developing new threat vectors is one of the capabilities that distinguishes threat actor groups.

The terms "threat vector" and "attack vector" are often taken to mean the same thing. Some sources distinguish the use of threat vector to refer to the analysis of the potential attack surface and attack vector to analyze an exploit that has been successfully executed.

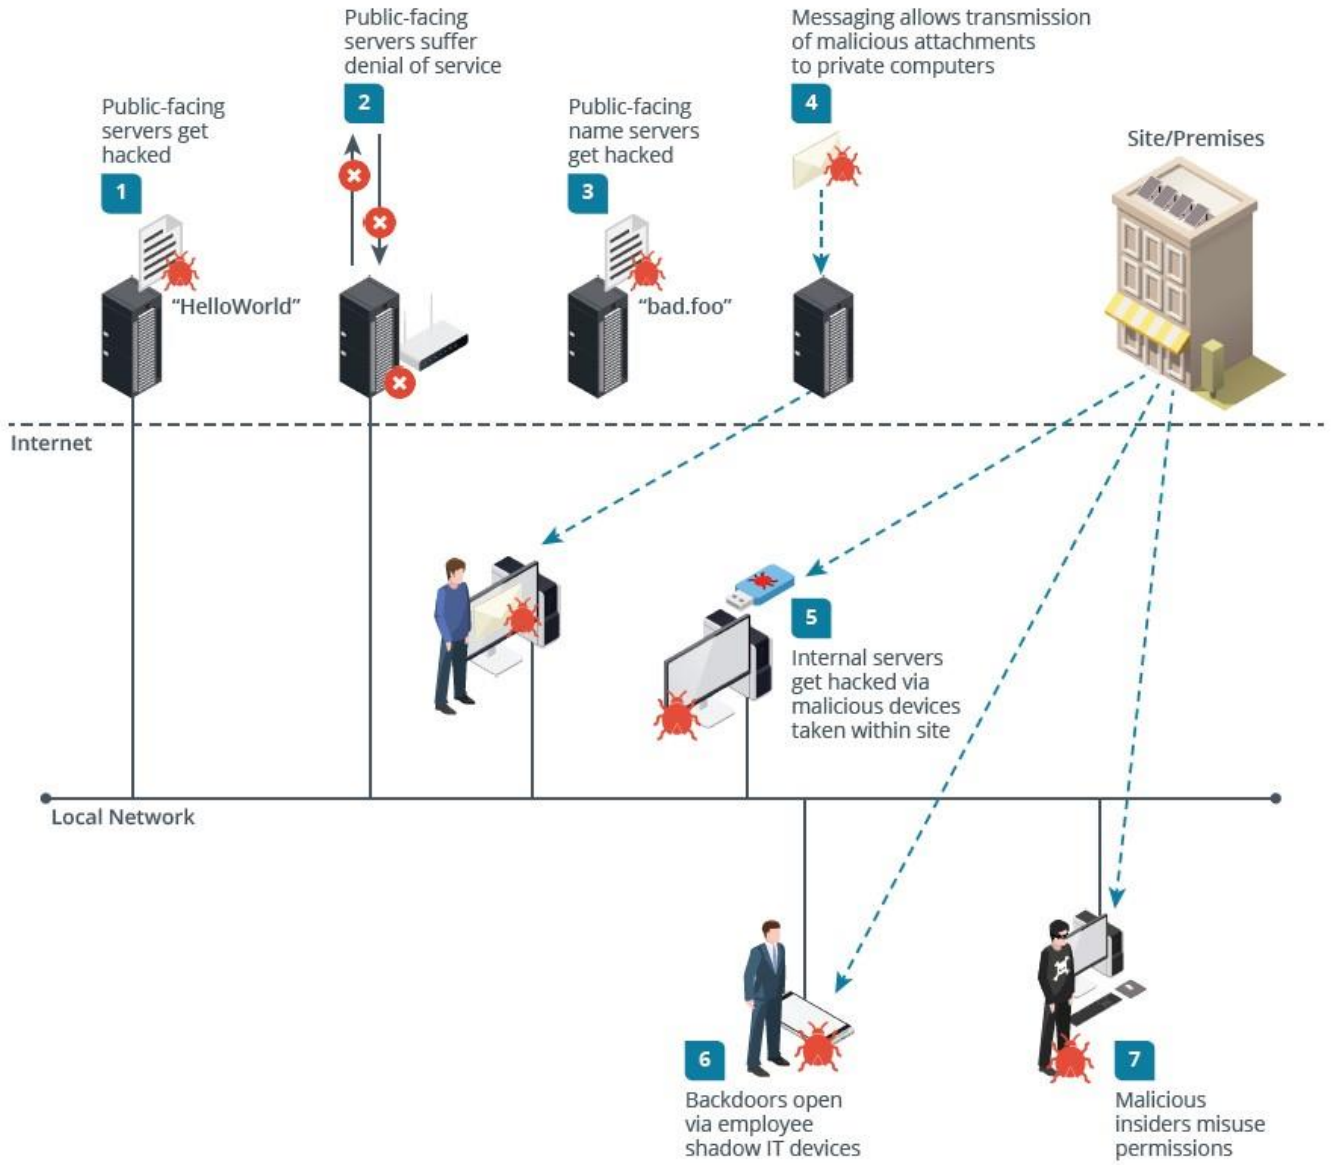

Description (Sample organization attack surface showing hacked servers, denial of service attacks, infected name servers, messaging with malicious attachments, malicious devices, shadow IT devices, and malicious insiders.)

The top row has four servers and a building,representing a company site, above a dotted line labeled as Internet. This indicates they are public facing servers and an organization.

The first two servers and four endpoints have a line connected to another line, below and parallel to the Internet line, labeled Local Network.

All of these devices represent a possible attack surface and there are seven vulnerabilities identified by labels next to the devices.

The first server represents public-facing servers that get hacked. It has a document with a bug icon and the word HelloWorld, refering to the virus by the same name.

The second server has a wireless router next to it and arrows pointing away and toward it. The arrows have red Xs on them indicating that traffic is stopped. It represents public-facing servers suffering denial of service (DoS) attack.

The third server represents public-facing name servers that are hacked. It also has the document and bug icons with the label bad.foo, refering to the DNS virus.

The fourth server has an envelope and bug icon and a dotted line to the first endpoint computer. This represents messaging that allows transmission of malicious attachments to private computers.

The second internal computer has a USB thumb drive near it and a dotted line to the public premises. It represents internal servers that are hacked via malicious devices within the site.

The third endpoint device is a smartphone. This represents backdoors opened via employee shadow IT devices.

Finally, the last endpoint device represents malicious insiders misuse of permissions.

Assessing the attack surface.

Vulnerable Software Vectors

Vulnerable software contains a flaw in its code or design that can be exploited to circumvent access control or to crash the process. Typically, vulnerabilities can only be exploited in specific circumstances and are often swiftly fixed—patched—by the vendor. However, because of the complexity of modern software and the speed with which new versions must be released to market, almost no software is free from vulnerabilities. Also, an organization might not have an effective patch management system. Consequently, vulnerable software is a commonly exploited threat vector.

A large number of operating systems and applications run on a company's appliances, servers, clients, and cloud networks directly increases the potential attack surface. This attack surface can be reduced by consolidating to fewer products and by ensuring the same version of a product is deployed across the organization.

The impact and consequences of a software vulnerability are varied. As two contrasting examples, consider vulnerabilities affecting Adobe's PDF document reader versus a vulnerability in the server software underpinning transport security. The former could give a threat actor a foothold on a corporate network via a workstation; the latter could compromise the cryptographic keys used to provide secure web services. Both are potentially high-impact for different reasons.

Unsupported Systems and Applications

Unsupported systems and applications are a particular reason that vulnerable software will be exposed as a threat vector. An unsupported system is one where its vendor no longer develops updates and patches. Unless the organization can patch the faulty code, these services and apps will be highly vulnerable to exploits.

One strategy for dealing with unsupported apps that cannot be replaced is to try to isolate them from other systems. The idea is to reduce opportunities for a threat actor to access the vulnerable app and run exploit code. Using isolation as a substitute for patch management is an example of a compensating control.

Client-Based vs. Agentless

Scanning software helps organizations automate the discovery and classification of software vulnerabilities. These tools can also be used by threat actors as part of reconnaissance against a target. This scanning software can be implemented as a

client-based agent. The agent runs as a scanning process installed on each host and reports to a management server. Alternatively, the vulnerability management product might use agentless techniques to scan a host without requiring any sort of installation. Agentless scanning is most likely to be used in threat actor reconnaissance.

Network Vectors

Vulnerable software gives a threat actor the opportunity to execute malicious code on a system. To do this, the threat actor must be able to run exploit code on the system or over a network to trigger the vulnerability. An exploit technique for any given software vulnerability can be classed as either remote or local:

- **Remote** means that the vulnerability can be exploited by sending code to the target over a network and does not depend on an authenticated session with the system to execute.
- **Local** means that the exploit code must be executed from an authenticated session on the computer. The attack could still occur over a network, but the threat actor must use some valid credentials or hijack an existing session to execute it.

Consequently, to minimize risks from software vulnerabilities, administrators must reduce the attack surface by eliminating unsecure networks. An unsecure network is one that *lacks* the attributes of confidentiality, integrity, and availability:

- **Lack of Confidentiality** —threat actors are able to snoop on network traffic and recover passwords or other sensitive information. These are also described as *eavesdropping attacks* .
- **Lack of Integrity** —threat actors can attach unauthorized devices. These could be used to snoop on traffic or intercept and modify it, run spoofed services and apps, or run exploit code against other network hosts. These are often described as *on-path attacks* .
- **Lack of Availability** —threat actors are able to perform service disruption attacks. These are also described as *denial of service (DoS) attacks* .

A secure network uses an access control framework and cryptographic solutions to identify, authenticate, authorize, and audit network users, hosts, and traffic.

Some specific threat vectors associated with unsecured networks are as follows:

Threat Vector	Description
Direct access	The threat actor uses physical access to the site to perpetrate an attack. Examples could include getting access to an unlocked workstation, using a boot disk to try to install malicious tools, or physically stealing a PC, laptop, or disk drive.
Wired network	A threat actor with access to the site attaches an unauthorized device to a physical network port, and the device is permitted to communicate with other hosts. This potentially allows the threat actor to launch eavesdropping, on-path, and DoS attacks.
Remote and wireless network	The attacker either obtains credentials for a remote access or wireless connection to the network or cracks the security protocols used for authentication. Alternatively, the attacker spoofs a trusted resource, such as an access point, and uses it to perform credential harvesting and then uses the stolen account details to access the network.
Cloud access	Many companies now run part or all of their network services via internet-accessible clouds. The attacker only needs to find one account, service, or host with weak credentials to gain access. The attacker is likely to target the accounts used to develop services in the cloud or manage

Threat Vector	Description
	cloud systems. They may also try to attack the cloud service provider (CSP) as a way of accessing the victim system.
Bluetooth network	The threat actor exploits a vulnerability or misconfiguration to transmit a malicious file to a user's device over the Bluetooth personal area wireless networking protocol.
Default credentials	The attacker gains control of a network device or app because it has been left configured with a default password. Default credentials are likely to be published in the product's setup documentation or are otherwise easy to discover.
Open service Port	The threat actor is able to establish an unauthenticated connection to a logical TCP or UDP network port. The server will run an application to process network traffic arriving over the port. The software might be vulnerable to exploit code or to service disruption.

Servers have to open necessary ports to make authorized network applications and services work. However, as part of reducing the attack surface, servers should not be configured to allow traffic on any unnecessary ports. Networks can use secure design principles, access control, firewalls, and intrusion detection to reduce the attack surface.

Lure-Based Vectors

A lure is something superficially attractive or interesting that causes its target to want it, even though it may be concealing something dangerous, like a hook. In cybersecurity terms, when the target opens the file bait, it delivers a malicious payload hook that will typically give the threat actor control over the system or perform service disruption.

If the threat actor cannot gain sufficient access to run a remote or local exploit directly, a lure might trick a user into facilitating the attack. The following media are commonly used as lures:

- **Removable Device** —the attacker conceals malware on a USB thumb drive or memory card and tries to trick employees into connecting the media to a PC, laptop, or smartphone. For some exploits, simply connecting the media may be sufficient to run the malware. More typically, the attacker may need the employee to open a file in a vulnerable application or run a setup program.

In a drop attack, the threat actor simply leaves infected USB sticks on office grounds, reception areas, or parking lots, expecting at least one employee to pick one up and plug it into a computer.

- **Executable File** —the threat actor conceals exploit code in a program file. One example is Trojan Horse malware. A Trojan is a program that seems to be something free and useful or fun, but it contains a process that will create backdoor access to the computer for the threat actor.
- **Document Files** —the threat actor conceals malicious code by embedding it in word processing and PDF format files. This can take advantage of scripting features or simply exploit a vulnerability in the document viewer or editor software.
- **Image Files** —the threat actor conceals exploit code within an image file that targets a vulnerability in browser or document editing software.

These vectors expose a large and diverse attack surface, from the USB and flash card readers installed on most computers to the software used to browse websites and view/edit documents. Reducing this attack surface requires effective endpoint security management, using controls such as vulnerability management, antivirus, program execution control, and intrusion detection.

Message-Based Vectors

When using a file-based lure, the threat actor needs a mechanism to deliver the file and a message that will trick a user into opening the file on their computer. Consequently, any features that allow direct messaging to network users must be considered as part of the potential attack surface:

Threat Vector	Description
Email	The attacker sends a malicious file attachment via email or any other communications system that allows attachments. The attacker needs to use social engineering techniques to persuade or trick the user into opening the attachment.
Short Message Service (SMS)	The file or a link to the file is sent to a mobile device using the text messaging handler built into smartphone firmware and a protocol called Signaling System 7 (SS7). SMS and the SS7 protocol are associated with numerous vulnerabilities. Additionally, an organization is unlikely to have any monitoring capability for SMS as it is operated by the handset or subscriber identity module (SIM) card provider.
Instant messaging (IM)	There are many replacements for SMS that run on Windows, Android, or iOS devices. These can support voice and video messaging, plus file attachments. Most of these services are secured using encryption and offer considerably more security than SMS, but they can still contain software vulnerabilities. The use of encryption can make it difficult for an organization to scan messages and attachments for threats.
Web and social media	Malware may be concealed in files attached to posts or presented as downloads. An attacker may compromise a site so that it automatically infects vulnerable browser software (a drive-by download). Social media may also be used more subtly, such as a disinformation campaign that persuades users to install a "must-have" app that is actually a Trojan.

The most powerful exploits are zero-click. Most file-based exploit code has to be deliberately opened by the user. Zero-click means that simply receiving an attachment or viewing an image on a webpage triggers the exploit.

Message-based vectors can also be exploited by a threat actor to persuade a user to reveal a password or weaken the security configuration using some type of pretext. This type of attack might be perpetrated simply by placing a voice call to the user.

Supply Chain Attack Surface

A **supply chain** is an end-to-end process of designing, manufacturing, and distributing goods and services to customers. Rather than attack the target directly, a threat actor may seek ways to infiltrate it via companies in its supply chain. One high-profile example of this is the Target data breach, which was made via credentials held by the company's building systems vendor.

The process of ensuring reliable sources of equipment and software is called procurement management. In procurement management, it is helpful to distinguish several types of relationships:

- **Supplier** —obtains products directly from a manufacturer to sell in bulk to other businesses. This type of trade is referred to as business-to-business (B2B).
- **Vendor** —obtains products from suppliers to sell to retail businesses (B2B) or directly to customers (B2C). A vendor might add some level of customization and direct support.

- **Business Partner** —implies a closer relationship where two companies share closely aligned goals and marketing opportunities.

For example, Microsoft is a major software manufacturer and vendor, but it is not feasible to establish direct relationships with all its potential customers. To expand its markets, it develops partner relationships with original equipment manufacturers (OEMs) and solutions partners. Microsoft operates a program of certification and training for its partners, which improves product support and security awareness.

Each supplier and vendor has its own supply chain. For example, a motherboard manufacturer and supplier will use companies to fabricate individual chip components. The supply chain extends to distribution, so delivery companies and couriers are part of it.

This supply chain breadth and complexity expose organizations to a huge attack surface. For example, for a computer motherboard to be trustworthy, the supply chain of chip manufacturer, firmware code developer, OEM reseller, courier delivery company, and administrative staff responsible for provisioning the computing device to the end user must all be trustworthy. Anyone with the time and resources to modify the computer's firmware could create backdoor access. The same is true for any computer or network hardware, software, or service.

Establishing a trusted supply chain for computer equipment and services essentially means denying malicious actors the time or resources to modify the assets supplied. For most businesses, the use of reputable vendors will represent the best practical effort at securing the supply chain. Government, military/security services, and large enterprises will exercise greater scrutiny. Particular care should be taken if using secondhand machines.

The IT industry also depends on trade in services as well as physical assets. A managed services provider (MSP) provisions and supports IT resources such as networks, security, or web infrastructure. MSPs are useful when an organization finds it cheaper or more reliable to outsource all or part of IT provision rather than try to manage it directly. From a security point of view, this type of outsourcing is complex, as it can be difficult to monitor the MSP. The MSP's employees are all potential sources of insider threat.

2.1.8 Practice Questions (Section Quiz)

q_types_attackers_actor_secp8

Which of the following does a threat actor need in order to support a high-level sophisticated attack?

Answers:

- ***Resources and funding**
- System permissions
- Disinformation
- Data exfiltration

Explanation:

A high level of capability must be supported by resources and funding. Sophisticated threat actor groups need to be able to acquire resources, such as customized attack tools and skilled strategists, designers, coders, hackers, and social engineers. The most capable threat actor groups receive funding from nation-states and organized crime.

An internal/insider threat actor (such as an employee or contractor) has been granted permissions on the system. While permissions might be useful, if a threat actor requires resources and funding, the threat actor is normally external.

Disinformation and data exfiltration are strategies often used by an external threat actor. Data exfiltration transfers a copy of some type of valuable information from a computer or network without authorization. Disinformation falsifies some type of

trusted resource, such as changing the content of a website, manipulating search engines to inject fake sites, or using bots to post false information to social media sites.

q_types_attackers_advanced_secp8

A threat actor successfully breached an advanced corporate network, bypassing multi-factor authentication and intricate intrusion detection systems. The highly coordinated attack leveraged zero-day vulnerabilities and sophisticated custom-made malware.

Which of the following BEST describes the capability level of this threat actor?

Answers:

- ***Advanced**
- Novice
- Intermediate
- Unskilled

Explanation:

The description of the threat actor's actions suggests a high degree of sophistication and technical capability, usually associated with advanced threat actors.

A novice threat actor is typically someone with limited knowledge and hacking skills. The novice threat actor is unlikely to carry out a highly sophisticated attack on an advanced corporate network.

An intermediate threat actor might have some ability to bypass basic security controls. Still, it would be difficult for the actor to carry out a highly sophisticated attack leveraging zero-day vulnerabilities and custom-made malware.

An unskilled threat actor would have limited knowledge and capability to perform sophisticated attacks. The unskilled actor is unlikely to be able to breach an advanced corporate network.

q_types_attackers_apt_secp8

A nation-state developed cyber weapons to achieve military influence and has the ability to obtain and maintain access to compromised networks.

What should a cybersecurity team address when this occurs to their organization's compromised networks?

Answers:

- ***Advanced persistent threats (APT)**
- A hacktivist
- A hacker
- An insider threat

Explanation:

APTs are cyber nation-state adversaries that have developed cybersecurity expertise and use cyber weapons to compromise network security and achieve military and commercial goals.

Hacktivists use cyber weapons to promote a political agenda. They can attempt to obtain and release confidential information to the public domain, perform denial-of-service (DoS) attacks, or deface websites.

Hackers have the skills to gain access to computer systems through unauthorized or unapproved means. The term is sometimes associated with illegal or malicious system intrusion.

Insider threats are employees who harbor grievances or perpetrate fraud. For example, an insider threat might plan and execute a campaign to modify invoices and divert funds.

q_types_attackers_blackmail_secp8

A healthcare provider suddenly receives a threat from an unknown source claiming to have obtained sensitive patient data. The anonymous actor demands a significant sum of Bitcoin, threatening to release the information publicly if the provider does not make payment.

This kind of scenario BEST exemplifies which threat motivation?

Answers:

- ***Blackmail**
- Espionage
- Service disruption
- Disinformation

Explanation:

Blackmail involves threats to reveal damaging information unless the affected party meets the blackmailer's demands, such as demanding Bitcoin for not releasing stolen patient data.

Espionage typically refers to the covert practice of obtaining confidential information for strategic advantage, often conducted by nation-states or competitive companies.

Service disruption prevents an organization from functioning normally through distributed denial-of-service (DDoS) attacks or other means of halting services. The threat scenario mentioned does not include any attempt to halt the healthcare provider's services.

Disinformation refers to the dissemination of false information to deceive or mislead. The threat scenario mentioned does not involve spreading false information but rather the threat of exposing real, sensitive information.

q_types_attackers_data_exfiltration_01_secp8

What type of attack takes content from a local system, encrypts it, and sends it to the attacker's server via HTTP over the port 80?

Answers:

- ***Data exfiltration**
- DDoS attack
- Denial-of-service attack
- Input validation

Explanation:

Data exfiltration is an unauthorized copying or retrieval of data from a system. Data exfiltration attacks are a primary means for attackers to retrieve valuable data often destined for later sale on the black market.

Distributed denial-of-service (DDoS) is an attack that uses multiple compromised computers to launch the attack.

A denial-of-service (DoS) attack causes a service at a given host to fail or to become unavailable to legitimate users.

An input validation attack passes invalid data to the application. The input handling on the routine is inadequate, which causes the application or even the operating system (OS) to behave unexpectedly.

q_types_attackers_data_exfiltration_02_secp8

A threat actor can infiltrate an organization's network and silently extract sensitive proprietary data without detection. The data has a high value on the black market.

Which motivations BEST align with this threat actor's likely objective?

Answers:

- ***Data exfiltration**
- Disinformation
- Service Disruption
- Revenge

Explanation:

Data exfiltration transfers a copy of valuable information from a computer or network without authorization. Threat actors might perform this type of theft because they want the data asset for themselves, can exploit its loss as blackmail, or sell it to a third party.

Disinformation involves falsifying trusted resources, such as changing a website's content or manipulating search engines to inject fake sites.

Service disruption prevents an organization from working as it normally does. It involves an attack on a website or malware-blocking access to servers and employee workstations.

Revenge might involve service disruption or disinformation attacks, but there is no indication in the scenario that the threat actor's motivation was revenge.

q_types_attackers_espionage_secp8

A multinational corporation recently fell victim to a series of cyberattacks, disrupting services and leading to significant financial losses. After an investigation, the corporation found that these attacks were part of a systematic campaign to undermine the corporation's market position.

The highly sophisticated attacks suggest the involvement of a well-resourced entity with specific strategic objectives.

Which of the following motivations BEST describes this scenario?

Answers:

- ***Espionage**

- Revenge
- Financial
- Chaotic

Explanation:

Espionage, characterized by stealthy, long-term breaches, aims at acquiring secret information, often for strategic advantage. The intruders' focus on the proprietary designs and their ability to remain undetected aligns with this motivation.

Revenge motivations typically involve a disgruntled individual seeking retaliation. This scenario does not provide evidence of a personal grievance or individual retaliation.

While financial motivation often involves monetary gain through methods such as blackmail, extortion, or fraud, the primary goal in this scenario is acquiring proprietary information, not explicit financial gain.

Chaotic motivations aim to disrupt for its own sake, often as an act of vandalism or to sow chaos. However, this scenario's targeted, stealthy, and organized nature of the attacks suggests a more sophisticated motivation.

q_types_attackers_financial_gain_secp8

In a recent incident, a hacker group infiltrated a global financial institution's systems and stole the credit card information of millions of customers. The valuable information was soon available on the dark web.

Based on the scenario, what is the MOST likely motivation of the hacker group?

Answers:

- ***Financial gain**
- Philosophical beliefs
- Service disruption
- Ethical concerns

Explanation:

Financial gain is a key motivator for many threat actors. In this scenario, the hacker group's main goal was to profit from the sale of stolen credit card information, which indicates financial gain as the primary motivation.

Philosophical beliefs as a motivation typically involve a group disrupting services or stealing data to pursue their ethical or philosophical beliefs. There is no indication in the scenario that the hacker group's motivation was philosophical.

Service disruption prevents an organization from operating normally, such as by launching a distributed denial of service (DDoS) attack or using malware to block system access.

Ethical concerns are typically associated with whistleblowers or insiders who act out of concern over an organization's unethical behavior.

q_types_attackers_hacker_secp8

Which of the following is the BEST definition of the term *hacker?*

Answers:

- ***A general term used to describe any individual who uses their technical knowledge to gain unauthorized access to an organization.**
- Any individual whose attacks are politically motivated.
- The most organized, well-funded, and dangerous type of threat actor.
- A threat actor whose main goal is financial gain.
- A threat actor who lacks skills and sophistication but wants to impress their friends or garner attention.

Explanation:

The term *hacker* is a general term used to describe any individual who uses their technical knowledge to gain unauthorized access to an organization.

The following are specific types of hackers, also known as threat actors:

- A hacktivist is any individual whose attacks are politically motivated.
- A nation state is the most organized, well-funded, and dangerous type of threat actor.
- An organized crime threat actor is a group of cybercriminals whose main goal is financial gain.
- A script kiddie is a threat actor who lacks skills and sophistication but wants to impress their friends or garner attention. Script kiddies carry out an attack by using scripts or programs written by more advanced hackers.

q_types_attackers_hacktivist_01_secp8

Which of the following threat actors seeks to defame, shed light on, or cripple an organization or government?

Answers:

- ***Hacktivist**
- Script kiddie
- Insider
- Competitor
- Nation state

Explanation:

A hacktivist is any individual whose attacks are politically motivated. Instead of seeking financial gain, hacktivists want to defame, shed light on, or cripple an organization or government. Hacktivists often work alone. Occasionally, they create unified groups with like-minded hackers. For example, the website wikileaks.org is a repository of leaked government secrets, some of which have been obtain by hacktivists.

Script kiddies are usually motivated by the chance to impress their friends or garner attention in the hacking community. Insider threat actors can be motivated by negative feelings toward their employer, bribes from a competitor, or personal financial gain. Competitors could be motivated by financial gain, competitor defamation, or obtaining industry secrets.

There are two primary motives for nation state attacks,seeking to obtain sensitive information (such as government secrets) or seeking to cripple the target's network or infrastructure.

q_types_attackers_hacktivist_02_secp8

A prominent multinational corporation has experienced an unexpected spike in unauthorized network traffic aimed at its web servers. Upon investigation, the corporation discovered that the goal of this traffic was to disrupt its online services rather than gain unauthorized access or steal data.

The attack started shortly after the corporation made a controversial policy decision that sparked a public backlash.

Which type of threat actor is MOST likely responsible?

Answers:

- ***Hacktivist**
- Insider threat
- Individual hacker
- Nation-state

Explanation:

The objective of the attack, disruption of online services following a controversial policy decision, aligns with the typical motives of a hacktivist. Hacktivists often use their skills to promote a social or political cause, which is the most likely the reason for the attack.

Insider threats typically seek to exploit their access to sensitive data or systems for personal gain, which does not align with the motives described in this scenario.

While an individual hacker might have the technical skills to carry out such an attack, the motivations described in the scenario align more closely with a hacktivist.

Nation-state actors typically conduct attacks for strategic or geopolitical gain rather than as a direct response to corporate policy decisions.

q_types_attackers_hacktivist_03_secp8l

A globally recognized fast-food chain recently experienced a cyber attack. The attackers have not shown interest in stealing sensitive data or disrupting operations but have defaced the company's website with messages promoting animal rights and the ethical treatment of livestock.

Based on this information, which type of threat actor is MOST likely responsible for this attack?

Answers:

- ***Hacktivist**
- Insider threat
- Individual hacker
- Nation-state

Explanation:

Hacktivists use their skills to promote a social or political cause. The attackers' objective (in this case, defacing a website to promote animal rights) aligns with typical hacktivist motives.

Insider threats typically involve employees or other individuals with authorized access to systems and information. They usually target sensitive data or disrupt operations for personal gain or revenge, not social or political reasons.

Individual hackers might conduct these types of attacks, but the political and social motives of the attack align more closely with hacktivism.

Nation-state actors are typically motivated by geopolitical strategy, not social or political activism. They tend to target other nations' infrastructure, government systems, or large corporations to steal sensitive data or cause disruption.

q_types_attackers_hacktivist_04_secp8

An environmental advocacy group uses cyber weapons to put companies at risk and promote its agenda.

This scenario illustrates what type of threat actor?

Answers:

- ***Hacktivists**
- Advanced persistent threats (APTs)
- Insider threats
- Hackers

Explanation:

Hacktivists use cyber weapons to promote an agenda, steal confidential information, perform denial-of-service (DoS) attacks, or deface websites. Environmental and animal advocacy groups may target companies in various industries.

APTs are cyber nation-state adversaries that have developed cybersecurity expertise and use cyber weapons to compromise network security and achieve military and commercial goals.

Insider threats are employees who harbor grievances or perpetrate fraud. An insider threat might plan and execute a campaign to modify invoices and divert funds.

Hackers have the skills to gain access to computer systems through unauthorized or unapproved means. The term is sometimes associated with illegal or malicious system intrusion.

q_types_attackers_insider_01_secp8

The IT manager in your organization proposes taking steps to deflect a potential threat actor. The proposal includes the following:

- Create and follow onboarding and off-boarding procedures.
- Employ the principal of least privilege.
- Have appropriate physical security controls in place.

Which type of threat actor do these steps guard against?

Answers:

- ***Insider**
- Script kiddie
- Hacktivist
- Competitor

Explanation:

Because insiders are one of the most dangerous and overlooked threats to an organization, you need to take the appropriate steps to protect against them, such as requiring mandatory vacations, creating and following onboarding and off-boarding procedure, employing the principal of least privilege, and having appropriate physical security controls in place.

A script kiddie is an individual who carries out an attack by using scripts or programs written by more advanced hackers.

A hacktivist is any individual whose attacks are politically motivated.

A competitor threat actor carries out attacks on behalf of an organization and targets competing companies.

q_types_attackers_insider_02_secp8

What type of threat actor is an individual or group with authorized access to an organization's systems and data that can potentially misuse access for malicious purposes?

Answers:

- ***Insider threat**
- Hacktivist
- Nation-state
- Unskilled attacker

Explanation:

Insider threats are individuals or groups with authorized access to an organization's systems and data. They could potentially misuse their access for malicious purposes. Insider threats can include employees, contractors, or any other individuals granted internal access.

Hacktivists are politically motivated cyber attackers. They do not usually have authorized access to an organization's systems or data.

Governments support nation-state actors that engage in highly sophisticated cyberattacks. They typically do not have authorized access to an organization's systems.

Unskilled attackers often use readily available tools to perform cyberattacks and generally do not have authorized access to an organization's systems.

q_types_attackers_insider_03_secp8

What type of threat actor will attempt to exploit their authorized access within an organization for revenge or financial gain?

Answers:

- ***Insider threat**
- Hacktivist
- Nation-state
- Unskilled attacker

Explanation:

Insider threats are individuals or groups with authorized access to an organization's systems and data. They could potentially misuse their access for malicious purposes. Insider threats can include employees, contractors, or any other individuals granted internal access.

Hacktivists are politically motivated cyber attackers. They do not usually have authorized access to an organization's systems or data.

Governments support nation-state actors that engage in highly sophisticated cyberattacks. They typically do not have authorized access to an organization's systems.

Unskilled attackers often use readily available tools to perform cyberattacks and generally do not have authorized access to an organization's systems.

q_types_attackers_insider_04_secp8

An employee suspected of modifying company invoices diverted funds from a company account to their private bank account.

What kind of malicious actor type does this describe?

Answers:

- ***Insider threat**
- Competitor
- Unskilled attacker
- Hacktivist

Explanation:

Insider threats are employees who harbor grievances or perpetrate fraud. An insider threat might plan and execute a campaign to modify invoices and divert funds.

Competitor attacks aim at theft, disrupting a competitor's business, or damaging its reputation. Employees who recently changed companies might facilitate competitor attacks and bring an element of insider knowledge with them.

An unskilled attacker uses hacker tools without understanding how they work or having the ability to craft new attacks. Unskilled attacker attacks have no specific target or reasonable goal other than gaining attention.

Hacktivists use cyber weapons to promote a political agenda. They obtain and release confidential information to the public, perform denial-of-service (DoS) attacks, or deface websites.

q_types_attackers_internal_threat_secp8

A cyber security analyst notices an unusual amount of data transmitted from an employee's company computer to an unknown external IP address. The employee has all necessary permissions to access externally transferred sensitive data.

What type of threat actor is MOST likely responsible for this situation?

Answers:

- ***Internal threat actor**
- Unskilled attacker
- Hacktivist

- Nation-state

Explanation:

An internal threat actor is an individual with permissions on the system, typically an employee or contractor. The scenario describes an employee with legitimate access to the data transmitting it to an unknown external IP address, which suggests that the threat actor is internal.

While an unskilled attacker may have some involvement in this situation, the information provided does not support this conclusion.

A hacktivist is someone who uses hacking as a means of expressing or promoting political or social views.

While nation-state actors can also perpetrate cyberattacks, they are typically external threats that do not usually possess the necessary permissions for data access, as described in the scenario.

q_types_attackers_nation_state_01_secp8

An international financial institution recently discovered a persistent and sophisticated cyber attack. The scale and sophistication of the attack suggest that the threat actor has access to significant resources.

The nature of the attack indicates that the threat actor operates with an extended timeline and appears motivated by strategic advantage rather than immediate financial gain.

Which type of threat actor is MOST likely involved?

Answers:

- ***Nation-state**
- Insider threat
- Hacktivist
- Individual hacker

Explanation:

Nation-state actors have the advanced capabilities, significant resources, and strategic motivations to carry out the sophisticated, long-term attack the financial institution discovered.

An insider threat typically does not have the resources and the persistent motivation to carry out a strategic, long-term attack like the one the financial institution discovered.

While hacktivists are sometimes motivated to target a financial institution, they typically do not have the advanced capabilities or resources to carry out this type of sophisticated attack.

An individual hacker usually lacks the significant resources, extended timeline, and strategic motivations behind this attack.

q_types_attackers_nation_state_02_secp8

Which type of threat actor is MOST likely to engage in cyber espionage with strategic or political motivations?

Answers:

- ***Nation-state**
- Competitors
- Hacktivist
- Organized crime

Explanation:

Nation-state actors often have the support of governments. Their activities, including cyber espionage, are typically motivated by strategic or political reasons.

Competitors might engage in cyber espionage against their rivals but are generally motivated by business competitiveness, not strategic or political reasons.

Hacktivists use hacking to express or promote a political agenda, but they are not typically involved in state-level strategic or political cyber espionage.

Organized crime groups engage in cybercrime activities like financial fraud, blackmail, or extortion for profit. They are typically not involved in strategic or political cyber espionage.

q_types_attackers_opp_secp8

A hacker scans hundreds of IP addresses randomly on the internet until they find an exploitable target.

What kind of attack is this?

Answers:

- Targeted attack
- Insider attack
- Nation-state attack
- ***Opportunistic attack**

Explanation:

In this scenario, the hacker is looking for an easy target and doesn't care what they are attacking. This is considered an opportunistic attack.

If the hacker had been targeting a certain individual, company, organization, or nation, it would have been considered a targeted attack.

An insider attack is accomplished by a threat agent who has authorized access to an organization and either intentionally or unintentionally carries out an attack.

A nation-state attack is accomplished by a threat agent that is a sovereign state who may wage an all-out war on a target and have significant resources and money at their disposal.

q_types_attackers_organized_crime_secp8

Which type of threat actor is MOST likely to engage in cybercrime activities, such as financial fraud, blackmail, or extortion for profit, and often operates across the Internet from a different jurisdiction than their victims?

Answers:

- ***Organized crime**
- Hacktivist
- Nation-state
- Unskilled attacker

Explanation:

Organized crime groups are often involved in cybercrime activities like financial fraud, blackmail, or extortion, seeking any opportunity for profit. They can operate across the Internet from a different jurisdiction than their victims, increasing the complexity of prosecution.

Hacktivists are politically motivated cyber attackers. They do not usually have authorized access to an organization's systems or data.

Governments support nation-state actors that engage in highly sophisticated cyberattacks. They typically do not have authorized access to an organization's systems.

Unskilled attackers often use readily available tools to perform cyberattacks and generally do not have authorized access to an organization's systems.

q_types_attackers_political_01_secp8

A multinational corporation recently fell victim to a series of cyberattacks, disrupting services and leading to significant financial losses. After an investigation, the corporation found that these attacks were part of a systematic campaign to undermine the corporation's market position.

The highly sophisticated attacks suggest the involvement of a well-resourced entity with specific strategic objectives.

Which of the following motivations BEST describes this scenario?

Answers:

- ***Political**
- Financial
- Chaotic
- Revenge

Explanation:

Political motivations typically involve strategic objectives to bring about change or achieve specific goals, often at a societal or governance level. This scenario's systematic, strategic, and sophisticated attacks and the intent to undermine the corporation's market position suggest a political motivation.

While financial motivation often involves monetary gain through methods such as blackmail, extortion, or fraud, this scenario does not highlight any explicit financial gain for the threat actor.

Chaotic motivations generally aim to disrupt for their own sake, often as an act of vandalism or to sow chaos. However, this scenario's systematic and strategic nature of the attacks suggests a more sophisticated motivation.

Revenge motivations typically involve a disgruntled individual seeking retaliation.

q_types_attackers_political_02_secp8

A group of threat actors disrupts the online services of an oil company due to their disagreement with the company's environmental policies. They believe their actions can force the company to change its practices.

This type of threat actor is primarily driven by what kind of motivation?

Answers:

- ***Political/philosophical beliefs**
- Service disruption
- Espionage
- Financial gain

Explanation:

The group, motivated by their philosophical beliefs about environmental responsibility, use their actions to bring about change in line with these beliefs.

Service disruption is a method that threat actors can use to achieve their objectives. In this case, service disruption is the action taken, not the motivation behind it.

Espionage refers to the act of obtaining secret or confidential information without the permission of the holder of the information. In this scenario, there is no mention of the group attempting to obtain confidential information.

The scenario does not suggest that the group's motivation is financial gain. They are not seeking to profit from their actions but to effect change in the company's policies.

q_types_attackers_revenge_secp8

A group of people lost their jobs after their company filed for bankruptcy. These employees formed a closed hacktivist group to fashion a zero-day exploit targeting specific Windows operating systems (OS) on the company network.

They will use internal influences to get the exploit onto the network.

Which of the following factors will greatly influence the success of this attack? (Select two.)

Answers:

- ***Revenge for hardship**
- ***Former colleague assistance**
- Criminal gang threats
- Nation-state influence
- Political motivation

Explanation:

Personal hardship after losing a job can motivate carrying out a revenge attack and seeing it through to the end. Revenge is a common motivator for insider threats.

Former colleagues, who still work for the company, are a good resource to influence and insert a zero-day exploit onto the network. These colleagues (or insider threats) may want to help because they are sympathetic toward the employees who lost their jobs.

Criminal gangs or syndicates will seek any opportunity for criminal profit, but typical activities are financial fraud and extortion against individuals and companies.

Nation-state actors' goals are primarily espionage and strategic advantage over another country. Governments back these actors with virtually unlimited resources.

A political motivation means that the threat actor uses an attack to bring about some type of change in society or governance. This type of motivation does not fit the scenario of employees losing their jobs.

q_types_attackers_service_disruption_secp8

A major online retailer experiences a sudden halt in its services during the peak holiday shopping season. It traces the cause back to an orchestrated distributed denial-of-service (DDoS) attack, which overwhelmed the retailer's servers with traffic, making it impossible for legitimate users to access the site.

This attack BEST aligns with which type of threat motivation?

Answers:

- ***Service disruption**
- Financial
- Disinformation
- Espionage

Explanation:

Service disruption attacks aim to prevent an organization from operating normally. In this case, the distributed denial-of-service (DDoS) attack disrupted the retailer's services.

The effects of the attack might have financial implications for the retailer, but the primary goal was not financial gain for the attackers.

Disinformation involves falsifying trusted resources, such as altering website content or using bots to spread false information. In this scenario, no evidence suggests that the attackers spread false information.

Espionage involves stealthy, long-term breaches aimed at acquiring secret information, often for strategic advantage. The DDoS attack in this scenario does not aim at acquiring any secret information.

q_types_attackers_shadow_it_secp8

The IT department at a large corporation noticed an unfamiliar software application running on its network. Upon investigation, they discovered that a team in the marketing department started using a new cloud-based project management tool to improve their workflow efficiency.

The team did consult with the IT department before implementing this tool.

In the context of cybersecurity threats, what does this situation BEST exemplify?

Answers:

- ***Shadow IT**
- Careless password management

- Insider threat
- Nation-state

Explanation:

Shadow IT refers to hardware, software, and services used within an organization without explicit approval from the IT department. The scenario described, where the marketing team started using a new project management tool without consulting the IT department, is a classic example of Shadow IT.

While this can be a security risk, careless password management refers to using weak passwords or sharing passwords with others.

Insider threat is a potential risk posed by individuals with inside information about the company's security practices, data, and computer systems.

Nation-state activity usually involves cyberattacks conducted by a state or state-sponsored actors against another state or corporation to gather intelligence or disrupt operations.

q_types_attackers_sophistication_secp8

A targeted attack has a budget that can allocate physical and human resources to achieve its goals.

This type of attack contains what attribute?

Answers:

- ***Sophistication**
- Opportunistic
- Unskilled attacker
- Known threats

Explanation:

One must consider an adversaries' sophistication and level of resources and funding. A targeted attack might use highly sophisticated tools backed by a budget that can allocate physical and human resources.

Opportunistic attacks might launch without much sophistication or funding, simply by using tools widely available on the Internet.

An unskilled attacker uses hacker tools without necessarily understanding how they work or having the ability to craft new attacks. Unskilled attacker attacks might have no specific target or any reasonable goal other than gaining attention.

Known threats, such as viruses or rootkits, Trojans, botnets, distributed denial-of-service (DDoS), or specific software vulnerabilities, are relatively straightforward to identify. Organizations can scan for these types of threats with automated software.

q_types_attackers_state_actors_01_secp8

Which malicious actors are likely to show great interest in another country's energy infrastructure and have unlimited resources to carry out espionage attacks?

Answers:

- ***State actors**
- Shadow IT
- Unauthorized hackers
- Semi-authorized hackers

Explanation:

The primary goals of state actors are espionage and strategic advantage. These actors receive government backing, have virtually unlimited resources, and are known to be particular about another country's energy and health network systems.

Shadow IT is an unintentional insider threat where users purchase or introduce computer hardware or software to the workplace with the sanction of their IT department. They do this without a procurement or security analysis process.

Unauthorized hackers have malicious intent. These hackers have limited resources, especially when working alone.

Semi-authorized hackers seek out vulnerabilities in a product or network without seeking approval. After informing companies about vulnerabilities, they do not exploit them but seek voluntary compensation or bug bounty.

q_types_attackers_state_actors_02_secp8

Experts at a scientific facility suspect that operatives from another government entity planted malware and are spying on one of their top secret systems.

Which attacker type is likely responsible based on the attacker's location and likely goals?

Answers:

- ***State actors**
- Criminal syndicates
- Unskilled attackers
- Hacktivists

Explanation:

State actors are responsible for many attacks, particularly on energy and health network systems. They typically work at arm's length from the national government that sponsors and protects them, maintaining plausible deniability.

A criminal syndicate can operate across the internet from a different jurisdiction than its victim, increasing the complexity of prosecution. Syndicates will seek any opportunity for criminal profit, typically financial fraud.

A script kiddie is someone who uses hacker tools without necessarily understanding how they work or having the ability to craft new attacks.

Hacktivists might attempt to obtain and release confidential information to the public domain, perform denial-of-service (DoS) attacks, or deface websites.

q_types_attackers_unskilled_01_secp8

An organization's system alerting tool detects a series of unsuccessful attempts of someone trying to gain unauthorized access to its servers. These attempts lack sophistication and appear to be using publicly available hacking tools.

Which type of threat actor is MOST likely responsible for these attempts?

Answers:

- ***Unskilled attacker**
- Nation-state
- Insider threat
- Hacktivist

Explanation:

Unskilled attackers, often called script kiddies, typically use widely available hacking tools and lack the knowledge to mount sophisticated attacks. The haphazard, unsuccessful attempts described are characteristic of this type of threat actor.

Nation-state attackers are highly sophisticated and typically involved in advanced and well-coordinated cyberattacks, which is not consistent with the pattern of the attempts described.

Insider threats generally have some level of authorized access to the systems, and their attacks do not necessarily follow the pattern of random, unsuccessful attempts described.

Hacktivists typically have a specific political or social agenda, and their attacks often target specific entities or services to further their cause. The described attempts lack any such focus or purpose.

q_types_attackers_unskilled_02_secp8

Which type of threat actor is MOST likely to initiate random, unsophisticated cyberattacks, often utilizing readily available hacking tools without a clear understanding of how they work?

Answers:

- ***Unskilled attacker**
- Nation-state
- Insider threat
- Hacktivist

Explanation:

Unskilled attackers, also known as script kiddies, commonly use widely available hacking tools without fully understanding them. Their attacks are often random and lack sophistication.

Nation-state attackers usually conduct advanced and well-coordinated cyberattacks. They are unlikely to use randomly targeted and unsophisticated attacks.

Insider threats typically have some degree of authorized access to the systems. They do not need simple attack tools; they may misuse their authorized access.

Hacktivists usually have a specific political or social motive. Their attacks typically target particular entities or services to advance their cause. They are not known for random and unsophisticated cyberattacks.

q_types_attackers_war_secp8

A recent cyberattack led to massive disruptions in a country's power grid, causing widespread blackouts and significant economic and social damage. The country's cyber team traced the attack to a hostile nation-state's cyber warfare division.

In this case, what is the primary motivation of the perpetrators?

Answers:

- ***War**
- Financial gain
- Ethical concerns
- Levels of sophistication/capability

Explanation:

In this case, the hostile nation-state attacked to cause widespread disruption and damage, a common objective in warfare. Such acts are a part of the broader strategy of using cyber means to achieve military and political objectives.

Financial gain typically involves stealing data for selling on the dark web, extortion, or fraud.

Ethical concerns generally drive insiders or whistleblowers to act against an organization's unethical behavior. Ethical concerns do not apply to the hostile nation-state's actions, driven by strategic considerations rather than ethical issues.

While the attack demonstrates high sophistication and capability, it is not the primary motivation. Instead, it indicates the actor's resources and skills.

q_att_def_strat_attack_secp8

Match the general attack strategy on the left with the appropriate description on the right. (Each attack strategy may be used once, more than once, or not all.)

Answers:

- Reconnaissance
- Breaching
- Escalating privileges
- Staging
- Exploitation

Explanation:

General attack strategies include the following steps:

- Reconnaissance is the process of gathering information about an organization, including system hardware information, network configuration, and individual user information.
- A breach is the penetration of system defenses. Breaches are achieved using the information gathered during reconnaissance.
- An escalating privileges attack is one of the primary objectives of an attacker, which can be achieved by configuring additional (escalated) rights to do more than breach the system.
- Staging is preparing a computer to perform additional tasks in the attack, such as installing software designed to attack other systems.
- An Exploit is used to take advantage of known vulnerabilities in software and systems. Types of exploitation include stealing information, denying services, crashing systems, and modifying information.

q_att_def_strat_defense_secp8

Match the general defense methodology on the left with the appropriate description on the right. (Each methodology may be used once, more than once, or not all.)

Answers:

- Layering
- Principle of least privilege
- Variety
- Randomness
- Simplicity

Explanation:

General defense methodologies include the following items:

- Layering is the process of implementing multiple security measures to protect the same asset. Defense in depth or security in depth is the premise that no single layer is completely effective in securing the assets. The most secure system/network has many layers of security and eliminates single points of failure.
- When using the principle of least privilege, users or groups are given only the access they need to do their job and nothing more. When assigning privileges, be aware that it is often easier to give a user more access when they need it than to take away privileges that have already been granted.
- Defensive layers should have variety and be diverse. Implementing multiple layers of the exact same defense does not provide adequate strength against attacks.
- Randomness relies on the constant change in personal habits and passwords to prevent anticipated events and exploitation.
- Security measures should provide protection, but not be so complex that you do not understand and use them.

q_att_def_strat_privilege_secp8

Which of the following is the BEST example of the principle of least privilege?

Answers:

- Mary has been given access to all of the file servers.
- ***Wanda has been given access to the files that she needs for her job.**
- Jill has been given access to all of the files on one server.
- Lenny has been given access to files that he does not need for his job.

Explanation:

Wanda being given access only to what she needs to do her job is an example of the principle of least privilege.

The principle of least privilege states that users or groups are given only the access they need to do their jobs and nothing more.

q_att_def_strat_recon_secp8

In which phase of an attack does the attacker gather information about the target?

Answers:

- Escalating privileges
- ***Reconnaissance**
- Breach the system
- Exploit the system

Explanation:

Reconnaissance is the phase of an attack where the attacker is gathering information about the target. This can be done electronically using scanning tools or even physically by going through dumpsters.

Escalation of privileges comes at the end of the attack when the attacker gains access to unauthorized data.

Breaching or exploiting the system is when the attacker gains access to a system on the target network using a vulnerability.

q_att_def_strat_variety_secp8

You have decided to secure your system/network by including several layers of security.

What other defense methodology is recommended that you should you incorporate to these defense layers to enhance security?

Answers:

- ***Variety of methods**
- Simplicity
- Principle of least privilege
- Randomness

Explanation:

Besides securing your system network by including layers of security, the defensive layers should incorporate a variety of methods.

The following are other general defense methodologies that you should consider, but are not as closely tied to enhancing layering:

- Simplicity - Security measures should provide protection but not be so complex that it is difficult to understand and use.
- Principle of least privilege - The principle of least privilege states that users or groups are given only the access they need to do their jobs and nothing more.
- Randomness - Randomness in security is the constant change in personal habits and passwords to prevent predictable behavior.

q_attack_surfaces_bluetooth_secp8

What technique does the threat actor use in a Bluetooth network attack to transmit malicious files to a user's device?

Answers:

- ***Exploiting vulnerabilities or misconfigurations in the Bluetooth protocol**
- Spoofing a trusted access point to gain unauthorized access
- Obtaining credentials for remote access to the network

- Physically stealing a PC or laptop to execute the attack

Explanation:

In a Bluetooth network attack, the threat actor exploits vulnerabilities or misconfigurations in the Bluetooth protocol to transmit a malicious file to a user's device.

Spoofing a trusted access point to gain unauthorized access is not specific to Bluetooth network attacks but relates to the remote and wireless network threat vector.

Obtaining credentials for remote access to the network is also not specific to Bluetooth network attacks but pertains to the remote and wireless network threat vector.

Physically stealing a PC or laptop to execute the attack is part of the direct access threat vector and irrelevant to Bluetooth network attacks.

q_attack_surfaces_conceal_code_secp8

During a cybersecurity attack, how would a threat actor use image files as a lure to target a vulnerability in a browser or document editing software?

Answers:

- ***The threat actor conceals exploit code within an image file that targets a vulnerability in the browser or document editing software.**
- The threat actor conceals malware on a USB thumb drive or memory card and tricks employees into connecting the media to a PC, laptop, or smartphone.
- They may use a program file with concealed exploit code, like Trojan Horse malware, to create backdoor access.
- The threat actor embeds malicious code in word processing and PDF format files to exploit vulnerabilities in document viewer or editor software.

Explanation:

The image file simply serves as a lure. This lure entices the target to open it, thus delivering a malicious payload hook when executed and further triggering the exploit.

A removable device lure conceals malware on a removable device, which will trick employees into connecting it to their individual devices.

Concealing an exploit code in a program file like Trojan Horse malware is an executable file lure.

Embedding malicious code in word processing and PDF format files to exploit document viewer or editor software vulnerabilities is a document file lure.

q_attack_surfaces_direct_access_secp8

A threat actor gains physical access to an organization's premises and attempts to perpetrate an attack on the wired network.

What specific threat vector associated with unsecured networks is likely used by the threat actor in this scenario?

Answers:

- ***Direct access**
- Remote and wireless network
- Bluetooth network
- Default credentials

Explanation:

The threat actor gains physical access to the site, making the direct access threat vector the most relevant choice as it involves using physical access to perpetrate an attack, such as accessing an unlocked workstation or stealing a PC.

Remote and wireless network threat vector involves obtaining credentials for remote access or wireless connection to the network or spoofing a trusted resource like an access point.

Bluetooth network threat vector involves exploiting vulnerabilities or misconfigurations in Bluetooth devices, not gaining physical access to an organization's premises.

Default credentials threat vector refers to gaining control of a network device or app due to default passwords, not due to physical access to the organization's premises.

q_attack_surfaces_infiltrate_secp8

A hacktivist group is intercepting multiple emails between a company and a few vendors and has learned that the company is planning to purchase new laptops and some Universal Serial Bus (USB) thumb drives.

In what ways can the group breach the target company MOST effectively? (Select two.)

Answers:

- ***Infiltrate the shipping company**
- ***Add malicious USB drives**
- Create a fake social media account
- Steal an employee's laptop
- Obtain credentials for remote access to the network

Explanation:

Infiltrating the shipping company takes advantage of the supply chain. In this manner, malicious actors can replace parts of the laptop or hack the operating systems before it gets to the company.

Adding malicious USB thumb drives to the order is taking advantage of removable media to trick the user into plugging them into a computer where the hacker can carry out further attacks.

Creating a fake account on Facebook or LinkedIn takes advantage of social media to carry out social engineering attacks. Hackers can pose as the vendor in this case, but this is not as effective as directly infiltrating the company or adding malicious thumb drives to the order.

Stealing an employee's laptop uses direct access attacks involving physical or local breach techniques at the target person, system, or network. However, this option fails to input compromised assets back into company infrastructure.

Obtaining credentials for remote access to the network can contribute to the efforts, but, given the scenario, is not the most effective way to breach the target company.

q_attack_surfaces_instant_messaging_secp8

You are a cybersecurity analyst at a large organization. You have been tasked with identifying the most secure method of communication to mitigate the risk of message-based vectors.

Which of the following methods would be MOST effective?

Answers:

- Email
- Short Message Service (SMS)
- ***Instant messaging (IM)**
- Web and social media

Explanation:

Instant messaging (IM) is the correct answer. Many instant messaging platforms offer end-to-end encryption, which means that only the sender and the recipient can read the messages. This makes it more difficult for attackers to intercept and read the messages, making instant messaging the most secure method of communication in this list.

While email is a common method of communication in many organizations, it is also a common vector for cyber attacks. Phishing emails, for example, are a common method used by attackers to trick users into revealing sensitive information or downloading malware. Therefore, while email can be made more secure through various measures, it is not the most secure method of communication in this list.

Short Message Service (SMS), or text messaging, is another common method of communication. However, it is also susceptible to various forms of attack, including SMiShing (SMS phishing) and SIM swapping. Like email, while it can be made more secure, it is not the most secure method in this list.

While web and social media platforms are widely used for communication, they are also common targets for cyber attacks. Social engineering attacks, for example, often use social media to trick users into revealing sensitive information. Therefore, while these platforms can be made more secure, they are not the most secure method of communication in this list.

q_attack_surfaces_isolate_apps_secp8

CloudSecure is facing a cybersecurity challenge where some of its critical software applications are no longer supported by vendors, making them vulnerable to potential exploits. The IT team is exploring various strategies to mitigate the risk posed by these unsupported apps.

What is the MOST effective approach to enhance the security posture?

Answers:

- ***Isolating the unsupported apps from other systems to reduce the attack surface.**
- Implementing regular patch management to fix the faulty code.
- Consolidating all operating systems and applications into one product.
- Ignoring the vulnerability as it can only be exploited in specific circumstances.

Explanation:

Isolating the unsupported apps from other systems helps to prevent threat actors from accessing the vulnerable app and running exploit code, thus acting as a compensating control.

Unsupported apps do not receive updates and patches; therefore, regular patch management would not be possible.

Consolidating into one product may not address the central issue of unsupported apps, and it may not be feasible or even practical for the organization's diverse needs.

Ignoring the vulnerability is not recommended, as attackers can exploit it in specific circumstances, potentially causing additional harm to the organization.

q_attack_surfaces_lack_of_integrity_secp8

Which of the following involves threat actors attaching unauthorized devices to a physical network port, allowing them to eavesdrop on network traffic, intercept and modify data, run spoofed services and applications, or execute exploit code against other hosts?

Answers:

- ***Lack of integrity**
- Lack of availability
- Lack of confidentiality
- Lack of authentication

Explanation:

Commonly called on-path attacks, lack of integrity compromises the reliability of the data transmission.

A lack of availability is where the network is vulnerable to service disruption attacks, such as distributed denial-of-service (DDoS) attacks.

Lack of confidentiality results in the exposure of sensitive and private information during data transmission through the open port, making it susceptible to eavesdropping attacks by malicious actors.

A lack of authentication makes it viable for unauthorized users to access the network and its resources without proper credentials, increasing the likelihood of security breaches.

q_attack_surfaces_msp_secp8

CryptoCloud is expanding its business and is considering outsourcing its IT resources to a managed services provider (MSP) to improve efficiency and reliability.

Which of the following statements about MSPs and their role in the supply chain are correct? (Select two.)

Answers:

- ***MSPs may introduce a complex security challenge as monitoring their employees can be difficult.**
- ***MSPs primarily focus on providing support for IT resources such as networks, security, or web infrastructure.**
- MSPs handle the end-to-end process of designing, manufacturing, and distributing goods and services to customers.
- MSPs are only suitable for large enterprises with extensive IT infrastructure and are not recommended for smaller businesses.
- MSPs involves receiving an attachment or viewing an image on a webpage which triggers an exploit.

Explanation:

Outsourcing to an MSP can be complex from a security point of view due to the difficulty in monitoring the actions of the MSP's employees, who are potential sources of insider threats.

MSPs manage, monitor, and maintain the organization's IT infrastructure, applications, and services. Their services can include network management, security, data backup, cloud computing, hardware and software maintenance, help desk support, and other IT-related tasks.

MSPs are suitable for any organization of any size looking to outsource its IT resources.

The primary focus of MSPs is to provide IT resource support rather than handling supply chain logistics.

Zero-click means that simply receiving an attachment or viewing an image on a webpage triggers an exploit.

q_attack_surfaces_supply_chain_secp8

A group of hackers has been monitoring recent orders from a company involving new laptops and Universal Serial Bus (USB) thumb drives. The group infiltrated the shipping company and added malicious USB thumb drives to the order. The target company received the order without any concerns.

What vectors made this attack successful? (Select two.)

Answers:

- ***Supply chain**
- ***Removable media**
- Social media
- Direct access
- Cloud access

Explanation:

Infiltrating the shipping company takes advantage of the supply chain. Malicious actors can replace parts of the laptop or hack the operating systems before it gets to the company.

Adding malicious USB thumb drives to the order takes advantage of removable media to trick the user into plugging them into a computer where the hacker can carry out further attacks.

Creating a fake social media account posing as the vendor to carry out social engineering attacks is not as effective as infiltrating the company or adding malicious thumb drives to the order.

Stealing an employee's laptop uses direct access attacks involving physical or local breach techniques. However, this option fails to input compromised assets back into company infrastructure.

Cloud access involves an attacker finding one or more accounts, services, or hosts with weak credentials in the cloud to gain access to a network. However, in this scenario there is no need to hack the company's cloud.

q_attack_surfaces_vendor_secp8

As a cybersecurity analyst, you are tasked with reducing the supply chain attack surface in your organization.

Which of the following areas should you focus on to MOST effectively mitigate this risk?

Answers:

- Internal IT infrastructure
- Employee training
- ***Vendor management**
- Customer data protection

Explanation:

Vendor management is the correct answer. The supply chain attack surface is often increased by third-party vendors who have access to your organization's systems or data. By focusing on vendor management, including conducting regular security audits and enforcing strict security standards, you can significantly reduce the supply chain attack surface.

While securing the internal IT infrastructure is crucial for any organization, it does not directly address the supply chain attack surface. Supply chain attacks often come from external sources, such as third-party vendors, rather than internal systems.

While employee training is important for overall cybersecurity awareness, it does not directly address the supply chain attack surface. Employees can be trained to recognize and avoid common attack vectors, but this does not mitigate the risk of a compromised vendor.

While protecting customer data is a critical aspect of cybersecurity, it does not directly address the supply chain attack surface. Supply chain attacks often target the organization's systems or data, rather than customer data specifically.

2.2 Social Engineering

As you study this section, answer the following questions:

- What is social engineering?
- What are the phases of a social engineering attack?
- What is pretexting and how is it used in social engineering?
- What are some of the most common social engineering techniques?
- How are motivation techniques effective in convincing targets to comply with a hacker's desires?
- What are common variations of phishing?
- How does a watering hole attack work?

In this section, you will learn to:

- Use the Social Engineer Toolkit.
- Investigate a social engineering attack.
- Identify social engineering.

Key terms for this section include the following:

Term	Definition
Social engineering	An activity where the goal is to use deception and trickery to convince unsuspecting users to provide sensitive data or to violate security guidelines.

Impersonation	A social engineering attack where an attacker pretends to be someone they are not.
Pretexting	A social engineering tactic where a team communicates, whether directly or indirectly, a lie or half-truth in order to get someone to believe a falsehood.
Phishing	An email-based social engineering attack in which the attacker sends email from a supposedly reputable source, such as a bank, to try to elicit private information from the victim.
Vishing	A human-based attack where the attacker extracts information while speaking over the phone or leveraging IP-based voice messaging services (VoIP).
Smishing	A form of phishing that uses SMS text messages to trick a victim into revealing information.
Pharming	A type of attack that redirects users from a legitimate website to a malicious one.
Typosquatting	An attack in which an attacker registers a domain name with a common misspelling of an existing domain, so that a user who misspells a URL in a browser is taken to the attacker's website.
Business email compromise	An impersonation attack in which the attacker gains control of an employee's account and uses it to convince other employees to perform fraudulent actions.
Watering hole attack	An attack in which an attacker targets specific groups or organizations, discovers which websites they frequent, and injects malicious code into those sites.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.1 Compare and contrast common threat actors and motivations.
	• Threat actors
	  ○ Unskilled attacker
	  ○ Insider threat
	• Motivations
	
	2.2 Explain common threat vectors and attack surfaces.
	• Human vectors/social engineering
	  ○ Phishing
	  ○ Smishing
	  ○ Vishing

	○ Misinformation/disinformation ○ Impersonation ○ Business email compromise (BEC) ○ Pretexting ○ Watering hole ○ Brand impersonation ○ Typosquatting 5.6 Given a scenario, implement security awareness practices. • User guidance and training ○ Social engineering • Development
TestOut Security Pro	5.0 Audit and Security Assessment 5.2 Assessment Techniques 5.2.2 Identify Social Engineering

2.2.1 Social Engineering Overview (Lesson Video)

Transcript:

A social engineering attack involves human interaction. The attacker tries to trick their victim, convince them to do things, or make them give out information they wouldn't under normal circumstances.

Because this attack involves humans, it's very hard to protect your data and assets against social engineering. It's also much harder to track and catch it before any damage is done. The best way to minimize social engineer's impact is to educate employees, managers, administrators, custodial staff, clients, users, consumers--essentially, the entire planet should learn to recognize and respond appropriately to social engineering attacks.
Social engineers are master manipulators. They're very good at discerning signs and behaviors worth exploiting for information. Some of the tactics that exploit human weakness include moral obligation, innate human trust, threatening, asking for very little, and using ignorance.
If the attacker exploits the victim's willingness to be helpful and assist them out of a sense of responsibility, that's leveraging moral obligation.
If the attacker exploits the victim's natural tendency to trust others, they're taking advantage of innate human trust. The attacker wears certain clothes, creates a certain demeanor, and speaks words and terms the victim is familiar with so they'll trust the attacker and give them what they want.
A threat is when the attacker intimidates the victim by promising negative consequences if they don't comply with the attacker's request.
And attackers will often promise some reward in return for very little to nothing, asking the victim to do a small favor or share a bit of information in return for a huge payoff.
Ignorance means the victim isn't educated in social engineering tactics and prevention, so they can't recognize an attack even when it's happening to them.
Now let's talk about the steps involved in a social engineering attack. These steps can be divided into three main parts: research, development, and exploitation.
The research phase is when the attacker starts gathering information about the company or organization they'll attack. Attackers use a process called footprinting to do this.

This is when the attacker goes through the target company's official websites and social media, performs dumpster diving, searches for employees' names, email addresses, and IDs, goes on a company tour, and conducts other types of on-site observation.
Then he'll move on to the development phase, which has two parts: selecting individual targets within the company being attacked and forming a relationship with the targets.
Usually, attackers choose people who not only have access to the information or object they desire, but show signs of being frustrated, over-confident, or arrogant, or having other traits that make it easy to extract information from them.
Once he's selected his targets, the attacker will start forming a relationship with them through conversations, emails, shared interests, and so on. This builds the victims' trust in the attacker, allowing them to be comfortable and relaxed so that they'll be even more willing to help.
The exploitation phase is when the attacker takes advantage of the relationship with the victim and uses the victim to extract information, obtain access, or accomplish the attacker's purposes in some way. Some examples include disclosing a password and username over the phone; introducing the attacker to other company personnel, providing social proof for the attacker; inserting a USB flash drive with a malicious payload into a company computer; opening an infected email attachment; exposing trade secrets in a discussion with a supposed peer.
If the exploitation is successful, the only thing left to do is wrap things up without raising suspicion. Most attackers tie up loose ends, like erasing digital footprints and ensuring no items or information get left behind. A well-planned and smooth exit strategy is the attacker's goal and their final act in the exploitation phase.
Social engineering is such a prevalent threat because it involves humans. So, let's review the key things to remember about social engineers: skilled attackers will know which human weaknesses to exploit based on people's feelings, emotions, and personalities. They're also detail-oriented, patient observers. They're skilled manipulators and good at forming relationships. The best ones will even clean up their tracks to make sure they won't get caught even if the victim realizes what's happened some time later. And the best defense against social engineers is user education.

2.2.2 Social Engineering Overview Facts

Social engineering refers to an attacker enticing or manipulating people to perform tasks or relay information. Social engineering tries to get a person to do something the person would not do under normal circumstances.

This lesson covers the following topics:

- Social engineering
- Manipulation and motivation
- Social engineering process

Social Engineering

Social engineering refers to the means of either eliciting information from someone or getting them to perform some action for the threat actor. It can also be referred to as "hacking the human." A threat actor might use social engineering to gather intelligence as reconnaissance in preparation for an intrusion or to effect an actual intrusion by obtaining account credentials or persuading the target to run malware. There are many diverse social engineering strategies, but to illustrate the concept, consider the following scenarios:

- A threat actor creates an executable file that prompts a network user for their password and then records whatever the user inputs. The attacker then emails the executable file to the user with the story that the user must open the file and log on to the network again to clear up some login problems the organization has been experiencing that morning. After the user complies, the attacker now has access to their network credentials.
- A threat actor contacts the help desk pretending to be a remote sales representative who needs assistance setting up remote access. Through a series of phone calls, the attacker obtains the name/address of the remote access server and login credentials, in addition to phone numbers for remote access and for accessing the organization's private phone and voice-mail system.
- A threat actor triggers a fire alarm, slips into the building during the confusion, and attaches a monitoring device to a network port.

Manipulation and Motivation

Social engineers are master manipulators and use many techniques to motivate their targets to disclose information. The following table provides common examples of these techniques:

Technique	Description
Moral obligation	An attacker uses moral obligation and a sense of responsibility to exploit the target's willingness to be helpful.
Innate human trust	Attackers often exploit a target's natural tendency to trust others. The attacker wears the right clothes, has the right demeanor, and speaks words and terms the target is familiar with so that the target will comply with requests out of trust.
Threatening	An attacker may try to intimidate a target with threats to make the target comply with a request. This is especially the case when moral obligation and innate human trust tactics are ineffective.
Offering something for very little to nothing	Offering something for very little to nothing refers to an attacker promising huge rewards if the target is willing to do a very small favor. The small favor can include sharing what the target thinks is a very trivial piece of information for something the attacker offers.
Ignorance	Ignorance means the target is not educated in social engineering tactics and prevention, so the target does not recognize social engineering when it is happening. The attacker knows this and exploits the ignorance.
Authority and fear	Authority techniques rely on power to get a target to comply without questioning the attacker. The attacker pretends to be a superior with enough power that the target will comply right away without question. The attacker could also pretend to be there in the name of or upon the request of a superior. Authority is often combined with fear. If an authority figure threatens a target with being fired or demoted, the target is more likely to comply without a second thought.
Social proof	With a social proof technique, the attacker uses social pressure to convince the target that it is okay to share or do something. In this case, the attacker might say, "If everybody is doing it, then it is okay for you to do it, too."
Scarcity	Scarcity appeals to the target's greed. If something is in short supply and will not be available, the target is more likely to fall for it.
Likeability	Likeability works well because humans tend to do more to please a person they like as opposed to a person they do not like.
Urgency	To create a sense of urgency, an attacker fabricates a scenario of distress to convince an individual that action is immediately necessary.
Common ground and shared interest	Common ground and shared interest work because sharing a hobby, life experience, or problem instantly builds a connection and starts forming trust between two parties.

Social Engineering Process

The social engineering process can be divided into three main phases: research, development, and exploitation. The following table describes each phase.

Phase	Description
Research	In the research phase, the attacker gathers information about the target organization. Attackers use a process called *footprinting* , which takes advantage of all resources available to gain information. Footprinting includes going through the target organization's official websites and social media, performing dumpster diving, searching sources for employees' names, email addresses, and IDs, going through a tour of the organization, and other kinds of onsite observation.
Development	The development phase involves two parts: • Selecting individual targets within the organization being attacked. • Forming a relationship with the selected targets. Usually, attackers select people who will not only have access to the desired information or object but also show signs of being frustrated, overconfident, arrogant, or somehow easy to extract information from. Once a target is selected, the attacker will start forming a relationship with the target through conversations, emails, shared interests, etc. The relationship helps build the target's trust in the attacker, allowing the target to be comfortable, relaxed, and more willing to help.
Exploitation	In the exploitation phase, the attacker takes advantage of the relationship with the target and uses the target to extract information, obtain access, or accomplish the attacker's purposes in some way. Some examples include: • Disclosing password and username • Introducing the attacker to other personnel, thus providing social credibility for the attacker • Inserting a USB flash drive with a malicious payload into an organization's computer • Opening an infected email attachment • Exposing trade secrets in a discussion. If the exploitation is successful, the only thing left to do is to wrap things up without raising suspicion. Most attackers tie up loose ends, such as erasing digital footprints and ensuring no items or information are left behind for the target to determine whether an attack has occurred or identify the attacker. A well-planned and smooth exit strategy is the attacker's goal and final act in the exploitation phase.

2.2.3 Social Engineering Techniques (Lesson Video)

Transcript:

Adversaries can use a diverse range of techniques to compromise a security system. A prerequisite of many types of attacks is to obtain information about the network and security system. This knowledge is not only stored on computer disks; it also exists in the minds of people. Employees, contractors, suppliers, and customers represent part of the attack surface of any organization. Collectively, they're part of the attack surface referred to as the human vector.

Social engineering refers to eliciting information from someone or getting them to perform some action for the threat actor. It can also be referred to as "hacking the human." We'll look at some common vectors that highlight how cybercriminals exploit human psychology and communication channels to gain unauthorized access, steal sensitive information, or other goals. Knowing these techniques will help you to develop policies and other security controls to mitigate these risks.

Phishing is a combination of social engineering and spoofing. A phishing message might try to convince the user to perform some action, such as installing disguised malware or allowing a remote access connection. Other types of phishing campaigns use a spoofed website set up to imitate a bank, e?commerce site, or some other web resource that should be trusted by the target. The attacker then emails users of the real website indicating their account must be updated or with a hoax alert or alarm, supplying a link to the spoofed site. When the user authenticates with the spoofed site, their login credentials are captured.

Another variation on phishing is called smishing which uses SMS, or text messages instead of email. It's more difficult to determine if the links are spoofed due to the reduced screen and functionality. However, direct messages to a single contact have a high chance of failure.

Next is vishing, short for "voice phishing," which uses phone calls or voice messages to manipulate individuals under the guise of a trusted entity. One high-profile example happened in 2021 when a branch manager received a call that he thought was from his director asking him to transfer $35 million for an acquisition. The voice on the phone sounded just like his director but was a computer-generated voice.

The next vector is impersonation, which simply means pretending to be someone else. It's one of the basic social engineering techniques. Impersonation is possible when the target can't easily verify the attacker's identity, such as over the phone or via an email message. The classic impersonation attack is for the social engineer to phone into a department, claim they have to adjust something on the user's system remotely, and then get the user to reveal their password.

To make the impersonation more credible, hackers will utilize pretexting. This attack vector involves creating a fabricated scenario (pretext) where they carefully craft a story with convincing or intimidating details. To be successful, they depend on obtaining privileged information about their target. Information that might seem innocuous—such as department employee lists, job titles, phone numbers, diaries, invoices, or purchase orders—can help an attacker penetrate an organization through impersonation.

Another technique that overlaps with impersonation and pretexting is brand impersonation. Brand impersonation occurs when attackers create fake websites, emails, or social media profiles that mimic a legitimate brand's online presence. These fake assets are used to deceive individuals into providing personal information or credentials.

Business email compromise (BEC) is a more sophisticated variation of phishing that also employs impersonation. BEC refers to a sophisticated campaign that targets a specific individual within a company, typically an executive. To perpetrate this type of high-stakes attack, the threat actor might first try to gain control of a legitimate mail account to send the phishing messages. The threat actor is likely to perform reconnaissance to obtain a detailed understanding of the target and the best psychological approach and pretexts to trick them.

Disinformation/misinformation tactics could be used to create fake social media posts or referrers (sites that link to the fake site) to boost search ranking. Disinformation refers to a purposeful motivation to deceive. Misinformation refers to repeating false claims or rumors without the intention to deceive. A disinformation campaign might attempt to get others to repeat and amplify the false facts it creates as misinformation.

A watering hole attack relies on a group of targets that use an unsecured third-party website. For example, a threat actor might discover that the staff running an international e-commerce site might use a local pizza delivery firm. An attacker can compromise the pizza delivery firm's website so that it runs exploit code on visitors. They may be able to infect the computers of the e-commerce company's employees and penetrate the e-commerce company's systems.

Typosquatting means that the threat actor registers a domain name very similar to a real one, such as exannple.com, hoping that users won't notice the difference and assume they're browsing a trusted site or receiving email from a known source. Another technique is to register a hijacked subdomain using the primary domain of a trusted cloud provider, such as onmicrosoft.com. If a phishing message appears to come from example.onmicrosoft.com, many users will be inclined to trust it.

Well, that's it for this lesson. In this lesson, we looked at various ways that cybercriminals exploit human vectors that play on human nature through social engineering. The best defense against these types of attacks is training. Individuals and organizations that understand these methods will be more vigilant and less likely to fall prey to the attackers.

2.2.4 Social Engineering Techniques Facts

Adversaries can use a diverse range of techniques to compromise a security system. A prerequisite of many types of attacks is to obtain information about the network and security system. This knowledge is not only stored on computer disks; it also exists in the minds of employees and contractors. A person with permissions on the system is a potential target of manipulative threat actor techniques known as social engineering.

This lesson covers the following topics:

- Human vectors
- Impersonation and pretexting
- Phishing and pharming
- Typosquatting
- Business email compromise
- Brand impersonation and disinformation
- Watering hole attack

Human Vectors

People—employees, contractors, suppliers, and customers—operating computers and accounts are a part of the attack surface of any organization. Collectively, they are referred to as the human vector. Cybercriminals will attempt to exploit human psychology and communication channels to gain unauthorized access, steal sensitive information, or spread malware. Being able to compare and contrast social engineering techniques will help you lead security awareness training and develop policies and other security controls to mitigate these risks.

Impersonation and Pretexting

Impersonation simply means pretending to be someone else. It is one of the basic social engineering techniques. Impersonation is possible when the target cannot easily verify the attacker's identity, such as over the phone or via an email message. A threat actor will typically use one of two approaches:

- **Persuasive/consensus/liking** —convince the target that the request is a natural one and would be impolite or somehow "odd" to refuse.
- **Coercion/threat/urgency** —intimidate the target with a bogus appeal to authority or penalty, such as getting fired or not acting quickly enough to prevent some dire outcome.

The classic impersonation attack is for the social engineer to phone into a department, claim they have to adjust something on the user's system remotely, and then get the user to reveal their password.

Do you really know who's on the other end of the line? (Image © 123RF.com.)

The use of a carefully crafted story with convincing or intimidating details is referred to as **pretexting** . Making a convincing impersonation to either charm or intimidate a target usually depends on the attacker obtaining privileged information about the organization. For example, when the attacker impersonates a member of the organization's IT support team, the attack will be more effective with the identity details of the person being impersonated and the target.

Some social engineering techniques are dedicated to obtaining this type of intelligence as a reconnaissance activity. As most companies are set up toward customer service rather than security, this information is typically quite easy to come by. Information that might seem innocuous—such as department employee lists, job titles, phone numbers, diaries, invoices, or purchase orders—can help an attacker penetrate an organization through impersonation.

Phishing and Pharming

Phishing is a combination of social engineering and spoofing. It persuades or tricks the target into interacting with a malicious resource disguised as a trusted one, traditionally using email as the vector. A phishing message might try to convince the user to perform some action, such as installing disguised malware or allowing a remote access connection by the attacker.

Other types of phishing campaigns use a spoof website set up to imitate a bank or e-commerce site, or some other web resource that should be trusted by the target. The attacker then emails users of the genuine website to inform them that their account must be updated or with a hoax alert or alarm, supplying a disguised link that actually leads to the spoofed site. When the user authenticates with the spoofed site, their login credentials are captured.

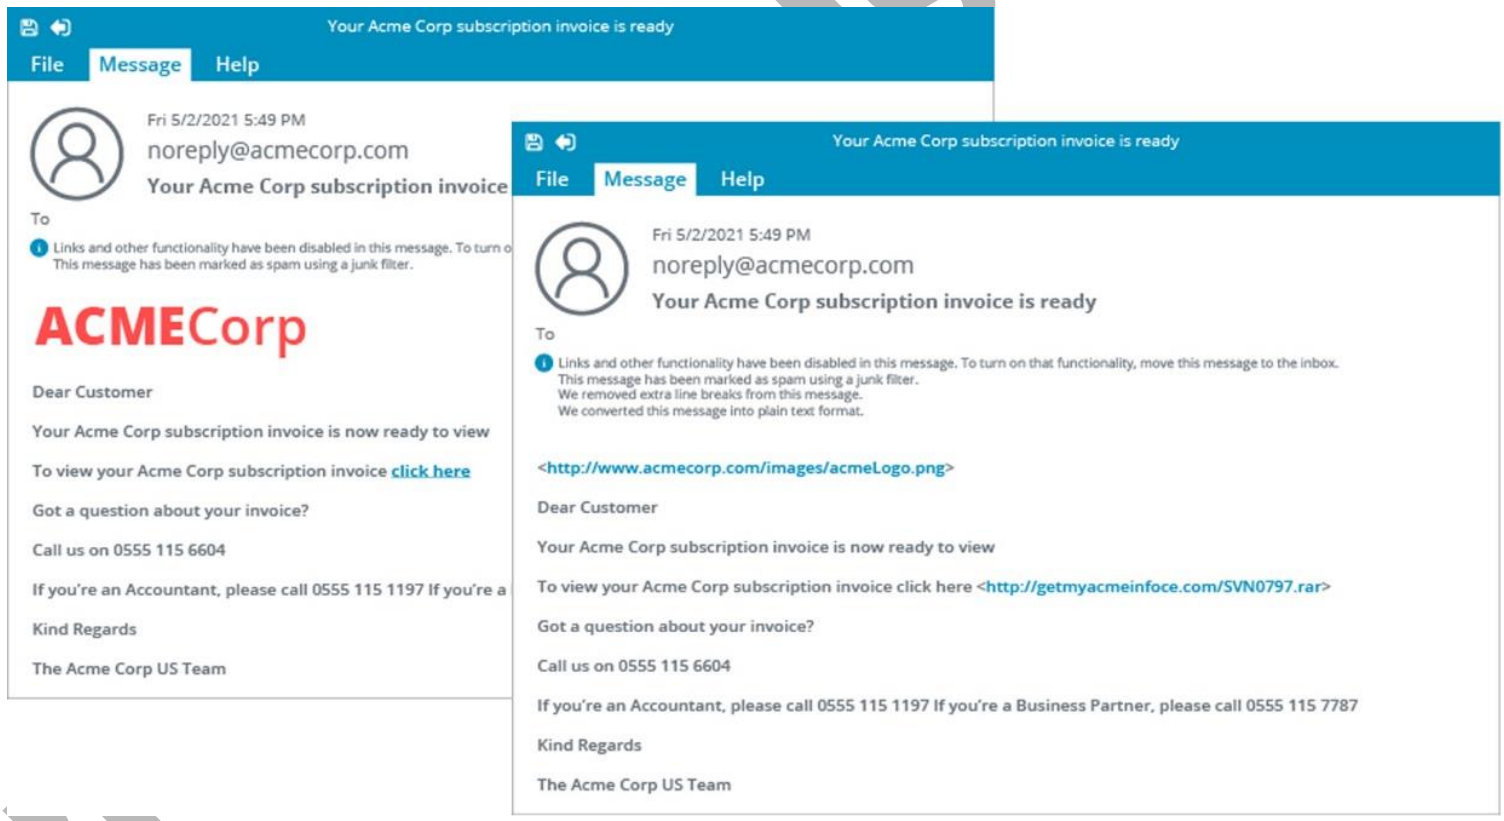

Example phishing email—On the right, you can see the message in its true form as the mail client has stripped out the formatting (shown on the left) designed to disguise the nature of the links.

Phishing refers specifically to email or text message threat vectors. The same sort of attack can be performed over other types of media:

- Vishing — a phishing attack conducted through a voice channel (telephone or VoIP, for instance). For example, targets could be called by someone purporting to represent their bank, asking them to verify a recent credit card transaction and requesting their security details. It can be much more difficult for someone to refuse a request made in a phone call than one made in an email.
- SMiShing — a phishing attack that uses Simple Message Service (SMS) text communications as the vector.

Rapid improvements in deep fake technology are likely to make phishing attempts via voice and even video messaging more prevalent in the future.

Direct messages to a single contact have a high chance of failure. Other social engineering techniques still use spoofed resources, such as fake sites and login pages, but rely on redirection or passive methods to entrap victims.

A pharming attack redirects users from a legitimate website to a malicious one. Rather than using social engineering techniques to trick the user, pharming relies on corrupting how the victim's computer performs internet name resolution so that they are redirected from the genuine site to the malicious one. For example, if mybank.foo should point to the IP address 2.2.2.2, a pharming attack would corrupt the name resolution process to make it point to IP address 6.6.6.6.

Typosquatting

Phishing and pharming both depend on impersonation to succeed. The spoofed message or site must appear to derive from a source that the target trusts. A threat actor can use various inconsistencies in how the message's source is represented in a mail client to trick the target into trusting the message source.

Email client software does not always identify the actual email address used to send the message. Instead, it displays a "From" field where a threat actor can add an arbitrary value. This technique is less common now as filtering software can be configured to alert the user to any discrepancy between the actual and claimed sender addresses.

Typosquatting means that the threat actor registers a domain name very similar to a real one, such as exannple.com , hoping that users will not notice the difference and assume they are browsing a trusted site or receiving email from a known source. These are also referred to as *cousin, lookalike,* or *doppelganger* domains.

Another technique is to register a hijacked subdomain using the primary domain of a trusted cloud provider, such as onmicrosoft.com . If a phishing message appears to come from example.onmicrosoft.com , many users will be inclined to trust it.

Business Email Compromise

Where phishing is typically associated with mass mailer attacks, **business email compromise** refers to a sophisticated campaign that targets a specific individual within a company, typically an executive or senior manager. The threat actor poses as a colleague, business partner, or vendor. The threat actor is likely to perform reconnaissance to obtain a detailed understanding of the target and the best psychological approach and pretexts to trick them. They are unlikely to use obvious features of mass mailer phishing messages such as spoofed links or malware file attachments.

To perpetrate this type of high-stakes attack, the threat actor might first try to gain control of a legitimate mail account to send the phishing messages.

Some sources use the term "business email compromise" to mean an attack with a specific financial motivation, where the objective is to persuade a budget holder to authorize a fraudulent payment or wire transfer. Similar terminology for highly targeted attacks includes spear phishing, whaling, CEO fraud (impersonating the CEO), and angler phishing (using social media as the vector).

Brand Impersonation and Disinformation

Brand impersonation means the threat actor commits resources to accurately duplicating a company's logos and formatting (fonts, colors, and heading/body paragraph styles) to make a phishing message or pharming website a visually compelling fake. The threat actor could even mimic the style or tone of email communications or website copy. They could try to get a phishing site listed high in search results using realistic content. **Disinformation/misinformation** tactics could be used to create fake social media posts or referrers (sites that link to the fake site) to boost search ranking.

Disinformation refers to a purposeful motivation to deceive. Misinformation refers to repeating false claims or rumors without the intention to deceive. A disinformation campaign might attempt to get others to repeat and amplify the false facts it creates as misinformation.

Watering Hole Attack

A watering hole attack relies on a group of targets that use an unsecured third-party website. For example, staff running an international e-commerce site might use a local pizza delivery firm. A threat actor might discover this fact through social engineering or other reconnaissance of the target. An attacker can compromise the pizza delivery firm's website so that it runs exploit code on visitors. They may be able to infect the computers of the e-commerce company's employees and penetrate the e-commerce company systems.

It is also important for everyone to be aware of what they post on their own social media platforms. That information can also be used to create more believable impersonations or attacks. As artificial intelligence, machine learning, or other technologies improve, they can be used to create what's known as deep fakes. Deep fakes create media that looks and or sounds like someone making statements that the person did not make or doing something in a picture or video that did not happen.

2.2.5 Use the Social Engineer Toolkit (Demo Video)

Transcript:

In this demo, we're going to talk about how to capture user credentials from a web page. While there are many different methods for gaining unauthorized access to a user's web account, we're going to do it by cloning a website's sign-in page and directing the user to enter the information on a fake page. To fool a user into going to our cloned sign-in page, we would create a phishing email and send it to them.

But in this demo, we're just going to cover cloning. To do this, we will use one of the many exploits provided by the Social Engineering Toolkit. Because this is a demo, we're going to use a fake website. I'll enter the credentials myself to show how the tool works.
First, we need to find a website that we would like to attack, But since we'll only be cloning one page of the site, we need to know what the full URL is for the login page. Let's go to www.fakewebsite.com. This is a Wordpress site, which makes our life a little easier, since the login page name is almost always the same. To make sure, though, we'll click the login button on the page and see where it takes us. I'll scroll down here and click Log in.
Now that we're on the login page, we can look in the address field and see the full URL of the exact page we want to clone. Since we'll need this URL later, we're going to copy it.
Now that we have the page we want to clone, let's open up the Social Engineering Toolkit. Since it's a command-line app, we'll be working in the terminal. Once the terminal opens, I'll type in 'setoolkit' and press Enter. On this screen, you can see that there are two attack options, Social Engineering Attacks and Penetration Testing. If you add extra attacks made by others, they'll appear under Third Party Modules. The remaining options are for maintaining the Social Engineering Toolkit. The option that we're looking for can be found under Social Engineering Attacks. So, I'll select number 1 and press Enter. We're interested in cloning a website, so we'll select option 2, Website Attack Vectors. To get to the clone option from here, we'll select number 3, Credential Harvester Attack Method, and then number 2, Site Cloner. Since we'll be using this machine to host the cloned page and listen for credentials, we can leave the default options. I'll just press Enter to continue.
Next, we'll use Ctrl+Shift+V to paste the URL we copied earlier and press Enter again. Now the Social Engineering Toolkit is saying, "Hey, do you understand what we're saying here?" Yes, yes I do. It also offers a couple of troubleshooting instructions in case the tool isn't working. Since we expect it to work, we'll just press Enter to continue.

Now we're finally at the page that will show us if the tool sees anything interesting. Earlier, I set up a website with a name that's just a little different from the name of the website we're cloning. This is the page that you would send to someone to try to trick them into opening it. It's basically the same address, but with the number 1 at the end. Let's go to that page in our web browser. I'm going to keep the terminal open on the side and open the web browser on the other side of the screen so you can see the tool working. I'll do that by opening Firefox and navigating to www.fakewebsite1.com. You can see that when I navigated to the cloned web page, something happened over to the right in the tool. It printed some information to the screen. This just shows that a connection request that was served to the requesting computer.
Now let's see what happens if I put in some user credentials. I'll type in 'admin' for the user and 'letmein' for the password. Press Enter to log in. And there we go. In the terminal, you can see some more information from the tool. Note that the tool tries to guess what information is important. This time, it correctly identified the password field, but didn't find a username field. That's because Wordpress username fields have kind of ambiguous names. But that didn't get in our way--this information was caught, as you can see right over here. It's worth noting that, after an attempted authentication, the Social Engineering Toolkit actually redirected us to the real website. This means that, typically, a user would likely just assume that they had entered their credentials incorrectly and try again. They probably wouldn't suspect a thing.
And that's it for this demo. We discussed how to use the Site Cloner in the Social Engineering Toolkit, and we saw how you--or anyone--can clone a website's look and feel and trick users into entering their credentials on a fake website.

2.2.6 Investigating a Social Engineering Attack (Demo Video)

Transcript:

In this demonstration, we'll talk about investigating a social engineering attack. We're going to focus on identifying the validity of emails, identifying the validity of websites, and identifying virus hoaxes.

Let's take a look at email validity. You should always be wary of emails from unknown organizations or senders. Be careful opening them, and be even more cautious with their attachments. A lot of hoax emails are pretty easy to spot because of their bad spelling and terrible grammar, but some of them look very, very real. Let's take a look at a couple of examples.
Let's start with this first one, Final Notice, Postal Status Notification. The first clue is the return address. It's from a variation of a Walmart address. It could be legit, but I'm not so sure. Next is the logo. It has a logo with a question mark after the word Walmart. I've never seen a Walmart logo like this, so I really wouldn't have to go any further to know it's a fake. If I hover over the logo, a bunch of addresses pop up. This looks like a batch of addresses that were spammed from this email. A little lower, when I hover over the Update Your Address Here link, I see a bunch more addresses. No legitimate email is going to be formatted to do this. Finally, here, it tells me to hurry, or I'll lose 17% of the money owed to me.
Let's look at this other one, the Check Your Account email. The address looks a bit odd. That's always the first clue. They want to confirm my information. I've never heard of this company, and I don't gamble, so within a second, I know this is fake. These links down here are Bitl.ly, which is a URL shortening service. I don't think most legitimate emails would use a shortening service like this.
Then we have another email that I took a screenshot of. It's this Attention email that says it's from usafed.gov. This is someone posing as somebody from the US Treasury. They're using a Gmail address for communication, which the government wouldn't do. At the bottom of the email, they've listed the legitimate website for the Department of the Treasury to make you think this is a legitimate email.
Another way you can verify an email is by taking a look at the header information. I have my Walmart email opened up in Outlook. This is another way to verify an email and its content. The header information is helpful because it shows IP information for the path an email took from its origin to the destination. It also shows other information, like the mail client that was used, the actual email address of the sender, and the originating IP address of the email.
There are a few ways to look at header information. In Outlook, it's helpful, but not easy to read. If we go into File > Properties, you can see the header information right here. As you can tell, this can't be resized, and it's hard to read in this format. The mail client in Windows 10 doesn't have a way to view header information, but don't worry. You can view header information by opening the email in a text editor like I've done with Notepad++.
By looking through the email in this format, you can pick out many things about its origin and determine whether it's legitimate. Remember, this email was supposed to be from Walmart. Here's the sender's originating IP. If we highlight this address and copy it, we can actually look up some information on that IP address to see where it's coming from.

All internet IP addresses for users come from some ISP somewhere. We can go to a website called arin.net, go to the search, and paste in our address. It'll find some information about this IP address. There are services that you can pay for that will give you a lot more information, but this one is free and simple. Looking at this information, I'm very skeptical that this email came from Walmart because the mailing address for this IP is from France.
Another thing we want to talk about is verifying websites. The first thing that you should always do is double-check the spelling of the website you're navigating to. If you're trying to go to PayPal and the address you end up at is paypals.com (with an s) or paypal.fakesight.com, then you're likely in the wrong place. You could accidentally give away your login information for PayPal (or a similar site) by entering it into a false site. You could also verify a site by viewing the SSL certification information. An SSL certificate (Secure Sockets Layer certificate) is a digital certificate that authenticates the identity of a website and encrypts the information sent to its server using SSL technology.
Let's look at two different browsers, Google Chrome and the Microsoft Edge browser. If we open both of these, there's a slight difference in how the SSL certificates are shown. In the Chrome browser, if you click on the lock in the address bar, it shows information about the certificate. I see that it's issued to google.com. I have a few tabs up here, Details and Certification Path. Under Certification Path, I see more information about the Certificate Authority and the certificate.
Now, if you go to Edge, open up Gmail, and click on the little lock, it'll show you the same information, but in a slightly different format. If there's no lock and no SSL certificate, you should be a little more cautious. Certificates aren't the focus of this lesson, but you want to make sure you know how to verify certificate information when you're not sure about a website.
Our last topic is identifying virus hoaxes. You've probably seen something like this before. Maybe one of your coworkers has come to you in a panic, worried about this type of error, and it's obviously very alarming. It's meant to be; this hoax is designed to scare or intimidate users into thinking they've contracted some type of virus. Normally, it says it's something like a Trojan, as you can see here. Often, there's a number you can call where they'll take payment over the phone for fixing the virus, but they never actually do anything.
There have been similar hoaxes that use the same tactics, so it's good to understand that viruses don't usually want to be discovered. Unless you're being notified by a trusted antivirus software that you or your administrator has installed, then it's probably just a hoax.
That's it for this lesson. We discussed verifying emails by looking at the HTML version and header information. We talked about verifying that a website is correct by checking the spelling of the address and the presence of a valid SSL certificate. We finished by discussing virus hoaxes.

2.2.7 Identify Social Engineering (Simulation)

Scenario

You work as the IT security administrator for a small corporate network in the United States of America. The company president has received several questionable emails that he is concerned may be malicious attacks on the company.

He has asked you to determine whether the emails are hazardous and to handle them accordingly.

In this lab, your task is to:

- Read each email and determine whether it is legitimate.
- Delete any emails that are attempts at social engineering.
- Keep emails that are safe.

Hold your mouse over the embedded links to see the actual URL in the status bar at the bottom of the screen.

Explanation

Complete this lab as follows:

1. From the Inbox of the WebEmail interface, highlight an email.
2. Read and explore the email and determine whether it is a legitimate email. This includes using your mouse to hover over suspicious attachments and links.
3. Take the appropriate action for each email:
    - If the email is an attempt at social engineering, from the menu bar, select **Delete** .
    - If the email is safe, do nothing.
4. Repeat steps 1 through 3 for each email. The following table lists the actions you should take for each email.

Email	Diagnosis	Action	Explanation for Action
Microsoft Windows Update Center New Service Pack	Phishing	Delete	This email has various spelling errors. The link does not direct you to a Microsoft website.
Joe Davis Re: Lunch Today?	Malicious Attachment	Delete	This email appears to be from a colleague; however, why would he fail to respond to your lunch question and send you a random attachment in return?
Executive Recruiting Executive Jobs	Whaling	Delete	Whaling uses tailored information to attack executives. Clicking the link could install malware that would capture sensitive company information. The link is pointing to a site in Germany (.de). It is suspicious that this organization would recruit executives from the USA.
Human Resources Ethics Video	Safe	Keep	While this email has an embedded link, it is digitally signed, as indicated by the green shield and checkmark. Therefore, you know it actually comes from your Human Resources department. When you hover over the link, you see that it is a secure link to the corporate web server.
Online Banking Department Payment Pending	Phishing	Delete	This is a carefully crafted attempt to get your bank account information. Hover over the link and notice that it does not direct you to your credit union website but to an unknown IP address. It is also very unlikely that a bank would delete your account for not verifying your information.

Grandma Jacklin FW: FW: FW: Virus Attack Warning	Hoax	Delete	Any email that asks you to forward it to everyone you know is probably a hoax. This email also contains very bad grammar.
Emily Smith Web Site Update	Spear Phishing	Delete	While this email appears to come from a colleague, notice that the link points to an executable file from a Russian domain name (.ru). A report file is more likely to have an extension of .pdf. .docx, .xlsx, or .txt. This probably is not a message a real colleague would send. This file will likely infect the computer with malware.
Sara Goodwin Wow!!	Malicious Attachment	Delete	Emails with attachments from unknown people who address you as "Dear Friend" are probably not safe.
Grandma Jacklin Free Airline Tickets	Hoax	Delete	Any email that asks you to forward it to everyone you know is probably a hoax, even if the contents promise you a prize. In addition, there is no way to know how many people the email has been forwarded to. Likewise, it is very unlikely that an airline would give away that many free tickets.
Human Resources IMPORTANT NOTICE-Action Required	Safe	Keep	While this email appears very urgent, it doesn't ask you to click on anything or run any attachments. It does inform you that you need to go to a website that you should already know and make sure your courses are complete.
Activities Committee Pumpkin Contest	Safe	Keep	This email doesn't ask you to click on anything or run any attachments.
Robert Williams Presentation	Safe	Keep	This email doesn't ask you to click on anything or run any attachments.

2.2.8 Social Engineering Techniques

2.2.9 Practice Questions (Section Quiz)

q_social_engr_ovw_authority_secp8

An organization's receptionist received a phone call from an individual claiming to be a partner in a high-level project and requesting sensitive information.

The individual is engaging in which type of social engineering?

Answers:

- Urgency
- Social proof
- Common ground
- ***Authority**

Explanation:

Authority social engineering entails an attacker either lying about having authority or using their high status in a company to force victims to perform actions that exceed their authorization level.

Urgency social engineering entails creating a sense of urgency, such as an attacker fabricating a scenario of distress to convince an individual that action is immediately necessary.

Social proof entails using social pressure to convince the target that it's okay to share or do something. For example, the attacker might say, "If everybody is doing it, then it's okay for you to do it, too."

Common ground social engineering entails establishing common ground and shared interest. This approach can be very effective because sharing a hobby, life experience, or problem instantly builds a connection and starts forming trust between two parties.

q_social_engr_ovw_dev_secp8

Ron, a hacker, wants to get access to a prestigious law firm he has been watching for a while. June, an administrative assistant at the law firm, is having lunch at the food court around the corner from her office.

Ron notices that June has a picture of a dog on her phone. He casually walks by and starts a conversation about dogs.

Which phase of the social engineering process is Ron in?

Answers:

- Research phase
- Exploitation phase
- ***Development phase**
- Elicitation phase

Explanation:

The development phase involves two parts. These are selecting individual targets within a company and forming a relationship with those individuals.

The exploitation phase is when the attacker takes advantage of the relationship with the victim and uses the victim to extract information, obtain access, or accomplish the attacker's purposes in some way.

The research phase is when the attacker starts gathering information about the target company or organization.

Elicitation is a technique used to extract information from a target without arousing suspicion.

q_social_engr_ovw_exploitation_secp8

You are a cybersecurity analyst at a large corporation. You notice that a particular employee has been receiving emails from an unknown sender who claims to be a new colleague from a different department.

The sender has been engaging in friendly conversation, asking about the employee's role, and subtly inquiring about certain company processes. Recently, the sender asked the employee to open an attachment that supposedly contains a funny meme.

What phase of the social engineering process does this scenario represent and what should be your immediate action?

Answers:

- Research phase - Inform the employee about the potential threat and advise them to stop communication.
- Development phase - Report the incident to the IT department for further investigation.
- ***Exploitation phase - Isolate the employee's system and conduct a thorough security scan.**
- None of the above - This is a normal interaction and no action is required.

Explanation:

The exploitation phase is the correct answer. The attacker is in the exploitation phase, where they are attempting to use the relationship they've built with the employee to achieve their goal, in this case, by trying to get the employee to open a potentially malicious attachment. The immediate action should be to isolate the employee's system and conduct a thorough security scan to check for any threats.

The research phase is incorrect because the research phase typically involves the attacker gathering information about the target organization, not directly interacting with an employee.

The development phase is incorrect because the development phase involves the attacker forming a relationship with the target, which has already occurred in this scenario. The attacker has moved beyond this phase by attempting to get the employee to open a potentially malicious attachment.

Taking no action is not correct. The scenario described is not a normal interaction and represents a potential security threat. It's important to take immediate action to prevent any potential harm to the organization's systems and data.

q_social_engr_ovw_intim_secp8

Social engineers are master manipulators.

Which of the following are tactics they might use?

Answers:

- ***Moral obligation, ignorance, and threatening**
- Shoulder surfing, eavesdropping, and keylogging
- Eavesdropping, ignorance, and threatening
- Keylogging, shoulder surfing, and moral obligation

Explanation:

Social engineers are master manipulators. Some of the most popular tactics they use are moral obligation, innate human trust, threatening, an easy reward, and ignorance.

Social engineering attacks include shoulder surfing, eavesdropping, USB and keyloggers, spam and spim, and hoaxes.

q_social_engr_ovw_princ_secp8

Any attack involving human interaction of some kind is referred to as what?

Answers:

- An opportunistic attack
- Attacker manipulation
- ***Social engineering**
- A white hat hacker

Explanation:

Social engineering refers to any attack involving human interaction of some kind. Attackers who use social engineering try to convince a victim to perform actions or give out information they wouldn't under normal circumstances.

An opportunistic attack is typically automated and involves scanning a wide range of systems for known vulnerabilities, such as old software, exposed ports, poorly secured networks, and default configurations.

A white hat hacker helps companies find vulnerabilities in their security infrastructure.

Social engineers are master manipulators and use multiple tactics on their victims.

q_social_engr_ovw_scarcity_secp8

As a security officer in a large corporation, you receive an email from an unknown sender claiming to be a high-ranking executive from another branch of your company. The email states that due to a sudden surge in demand for your company's product, there is a shortage of supply.

The sender requests immediate access to your inventory data to assess the situation and promises a significant bonus for your prompt cooperation. The sender also mentions that this is a confidential matter and should not be discussed with anyone else.

What would be the BEST response in this scenario?

Answers:

- Immediately provide the requested data to help the company in this crisis.
- ***Report the email to your supervisor and the IT department without responding to the sender.**
- Respond to the sender asking for more details about the situation.
- Ignore the email as it does not concern your department.

Explanation:

Reporting the email to your supervisor and the IT department without responding to the sender is the best response. The email has several red flags that suggest it could be a social engineering attack. The sender is unknown, the request is urgent, confidential, and involves access to sensitive data. The promise of a bonus is a form of the scarcity technique. It's important to report such incidents to your supervisor and IT department for further investigation.

It is not advisable to share sensitive data without verifying the identity of the requester, even if they claim to be from your company. The promise of a significant bonus and the urgency of the situation are tactics often used in social engineering attacks, specifically the scarcity technique.

Responding to the sender asking for more details about the situation is not advisable because engaging with the sender could potentially expose your system to further risk. It's best to report the incident to the appropriate parties within your organization.

Ignoring the email is not the best course of action because it leaves the potential threat unaddressed. It's important to report such incidents to help your organization stay aware of potential security threats.

q_social_engr_ovw_social_example_secp8

Which of the following is an example of a social engineering attack?

Answers:

- ***A call from a threat actor posing as a remote sales representative to obtain the login credentials to a remote access server from the help desk.**
- An employee sends information to HR, but an attacker secretly intercepts and manipulates the communication, unbeknownst to both employees.
- An attacker floods a website's server with fake requests, making it slow or unresponsive to legitimate users.
- A fake bank email is sent to recipients asking them to update their account info via a link that leads to a fake site, capturing login details.

Explanation:

The example illustrates a social engineering attack where the threat actor deceives the help desk into providing sensitive information through a persuasive phone call.

The scenario is a man-in-the-middle attack focused on both intercepting and relaying communications, which is not present in this scenario.

This answer choice is an example of a denial-of-service attack. This type of attack causes service disruptions.

This scenario is a phishing attack that uses a combination of social engineering and spoofing to persuade or trick the target into interacting with a malicious source disguised as a trusted one.

q_social_engr_ovw_social_proof_secp8

You decided to purchase a natural medication online based on testimonials from several customers.

Later you hear from a news report that the company selling the product has been indicted for fraud. As part of the charges, the testimonials were found to be fake.

Which of the following motivation techniques did the company use to entice you to purchase the natural medication?

Answers:

- ***Social proof**
- Authority
- Urgency
- Scarcity

Explanation:

With a social proof technique, the attacker uses social pressure to convince the target that it is okay to share or do something. For example, the attacker might say, "If everybody is doing it, then it's okay for you to do it, too." In this scenario, the social proof technique the company utilized was fake customer testimonials.

The following are additional motivational techniques that do not fit the scenario:

- Authority - Authority techniques rely on power to get a target to comply without questioning the attacker. The attacker pretends to be a superior with enough power that the target will comply right away without question.
- Urgency - To create a sense of urgency, an attacker fabricates a scenario of distress to convince an individual that action is immediately necessary.
- Scarcity - Scarcity appeals to the target's greed. If something is in short supply and will not be available, the target is more likely to fall for it.

q_social_engr_tqs_brand_impersonation_secp8

What social engineering technique involves the threat actor committing resources to accurately duplicate a company's logos, formatting, and communication style to make a phishing message or fake website visually compelling and convincing?

Answers:

- ***Brand impersonation**
- Pharming
- Vishing
- Spear phishing

Explanation:

Brand impersonation is a social engineering technique where the threat actor makes a significant effort to create a phishing message or a fake website that accurately duplicates a legitimate company's logos, formatting, and communication style. This level of detail makes the phishing attempt visually compelling and convincing to trick the target into believing it's legitimate.

Pharming involves redirecting users from legitimate websites to malicious ones by corrupting the victim's computer's name resolution process.

Vishing is a type of phishing attack conducted through a voice channel. It is not explicitly related to duplicating a company's branding to create a visually compelling phishing message.

Spear phishing is a highly targeted social engineering attack focusing on a specific individual within the company.

q_social_engr_tqs_business_email_secp8

The senior manager at CyberCorp receives an email from what appears to be a trusted colleague within the company. The email requests sensitive financial information to complete an urgent transaction and looks legitimate, displaying the colleague's name, company logo, and formatting.

What type of sophisticated phishing attack occurs in this scenario?

Answers:

- ***Business email compromise**
- Mass mailer phishing
- Angler phishing
- Whaling

Explanation:

Business email compromise (BEC) is a sophisticated attack that targets specific individuals, such as executives. The threat actor impersonates a trusted colleague, business partner, or vendor to trick the target into performing actions or disclosing information.

Mass mailer phishing is a traditional form of phishing where many emails sent to a broad audience deceive targets into clicking on spoofed links or downloading malware.

Angler phishing involves using social media as the vector to deceive targets into clicking on malicious links or disclosing information.

Whaling is a phishing attack that targets high-level executives or decision-makers in an organization. It is like business email compromise but may not involve posing as a colleague or business partner.

q_social_engr_tqs_deep_fake_secp8

Which of the following is the term used for creating media that looks and or sounds like someone making statements that the person did not make?

Answers:

- ***Deep fake**
- Impersonation
- Brand impersonation
- Disinformation

Explanation:

Deep fakes create media that looks and or sounds like someone making statements that the person did not make or doing something in a picture or video that did not happen.

Impersonation involves an attacker posing as someone with a legitimate need for information. For example, someone might pose as your banker and contacts you to tell you that you are a victim of identity theft, and she will be glad to help you reset your account password. Here the attacker is impersonating your banker and she is contacting you on the pretext of identity theft.

Brand impersonation occurs when attackers create fake websites, emails, or social media profiles that mimic a legitimate brand's online presence. These fake assets are used to deceive individuals into providing personal information or credentials.

Disinformation is the intentional presentation of false information with the intent to influence decisions and actions. Disinformation can be the result of:

- Removing important facts.
- Removing contextual information from articles, statements, videos, or photos.

q_social_engr_tqs_duplicate_logos_secp8

Which disinformation/misinformation tactics create convincing brand impersonation for phishing attacks or pharming websites? (Select two.)

Answers:

- ***By accurately duplicating a company's logos and formatting, such as fonts, colors, and styles, to make the fake site visually compelling.**
- ***By mimicking the style or tone of email communications or website copy to create a convincing fake.**
- Using disinformation to purposefully deceive individuals by spreading false claims and rumors with the intention of causing confusion and harm.
- Repeating false claims or rumors without the intention to deceive others but aiming to get those false facts amplified by others.
- Primarily focusing on providing support for IT resources such as networks, security, or web infrastructure.

Explanation:

A disinformation/misinformation campaign utilizes accurate duplications of a company's logos and formatting to appear more legitimate.

Using the mimicry of communication styles creates convincing brand impersonation. This mimicry can lead to successful phishing attacks or pharming websites that appear visually compelling and authentic.

While disinformation campaigns aim to purposely deceive individuals with false claims, it does not directly relate to brand impersonation or the creation of visually compelling fake messages or websites.

Repeating rumors, false claims, or misinformation does not address brand impersonation or the use of realistic content to boost search rankings for fake websites.

MSPs primarily focus on providing support for IT resources such as networks, security, or web infrastructure.

q_social_engr_tqs_impersonation_01_secp8

Impersonation is a social engineering attack that involves which of the following?

Answers:

- ***Pretending to be someone else and may use persuasive or coercive approaches to manipulate the target.**
- Manipulating the Domain Name System (DNS) to redirect website traffic.
- Sending deceptive emails to trick users into revealing sensitive information.
- Using malware to compromise a network and steal data.

Explanation:

Impersonation is a social engineering attack where the threat actor pretends to be a legitimate user or entity to gain unauthorized access, deceive others, or carry out fraudulent activities.

This option describes a pharming attack, not impersonation. Pharming manipulates Domain Name System (DNS) settings to redirect users to fraudulent websites.

This option refers to a phishing attack, not impersonation. Phishing uses deceptive emails to trick users into revealing sensitive information.

Unrelated to impersonation, this option describes a cyber attack involving malware. Malware attacks aim to compromise networks and steal data, not imitate legitimate users or entities.

q_social_engr_tqs_impersonation_02_secp8

Which social engineering technique involves pretending to be someone else and may use persuasive or coercive approaches to manipulate the target?

Answers:

- ***Impersonation**
- Pharming
- Phishing
- Watering hole attack

Explanation:

Impersonation involves pretending to be someone else and may use persuasive or coercive approaches to deceive the target.

Pharming is an attack that redirects users from a legitimate website to a malicious one but does not actually involve impersonation.

Phishing is a combination of social engineering and spoofing, aiming to trick the target into interacting with a malicious resource, but it doesn't necessarily involve impersonation.

A watering hole attack targets a group of users by compromising a third-party website they may commonly visit, but it is not related to impersonation.

q_social_engr_tqs_impersonation_03_secp8

A social engineer intercepted an end-user's phone call to an internet service provider (ISP) about a home internet outage.

Pretending to be the caller reporting the outage, the attacker immediately contacted the ISP to cancel the service call, dressed up as an internet tech, and then proceeded to enter the end-user's home with permission.

What social engineering attack did the ISP and end-user fall victim to?

Answers:

- ***Impersonation**
- Hoax
- Pharming
- Tailgating

Explanation:

Impersonation is a social engineering attack in which the attacker pretends to be someone else.

In a hoax attack, an email alert or web pop-up will claim to have identified a security problem, such as a virus infection, and offer a tool to fix the problem. The tool, of course, will be some sort of Trojan application.

Pharming relies on corrupting how the victim's computer performs internet name resolution to redirect them from the genuine site to the malicious one.

Tailgating is entering a secure area without authorization by following close behind the person with permission to open the door or checkpoint.

q_social_engr_tqs_phishing_01_secp8

Which of the following attacks tricks victims into providing confidential information (such as identity information or logon credentials) through emails or websites that impersonate an online entity that the victim trusts?

Answers:

- Pharming
- Pretexting
- ***Phishing**
- Preloading

Explanation:

Phishing tricks victims into providing confidential information, such as identity information or logon credentials, through emails or websites that impersonate an online entity that the victim trusts. These entities could include a financial institution or well known e-commerce site. Phishing is a specific form of social engineering.

Pharming involves the attacker executing malicious programs on the target's computer so that any URL traffic redirects to the attacker's malicious website.

Pretexting is conducting research and information gathering to create convincing identities, stories, and scenarios to be used on selected targets.

Preloading is used to set up a target by influencing the target's thoughts, opinions, and emotions.

q_social_engr_tqs_phishing_02_secp8

Customers receive a seemingly genuine email from their trusted bank informing them that their passwords need updating. However, when authenticating, an attacker captures the customer's credentials.

What kind of attack did the bank customers experience?

Answers:

- ***Phishing**
- Vishing
- SMiShing
- Whaling

Explanation:

Phishing is a combination of social engineering and spoofing, where the attacker sets up a spoof website to imitate a trusted one. The attacker then emails users of the genuine website, informing them that their accounts need updating and supplying a disguised link that leads to their spoofed site. Users then authenticate with the spoofed site, capturing their login credentials.

Vishing describes a phishing attack conducted through a voice channel (telephone or Voice over Internet Protocol, also called VoIP, for instance).

Smishing refers to fraudulent SMS texts. Other vectors could include instant messaging or social media sites.

A spear-phishing attack directed specifically against upper levels of management in the organization is sometimes called whaling.

q_social_engr_tqs_phishing_03_secp8

An employee receives an email from an unknown sender claiming to be from the IT department.

The email states that there is a login issue on the network and that the user needs to run the file to resolve the problem. The executable file prompts the user to input a network password, which the threat actor records.

What social engineering technique is the threat actor using in this scenario?

Answers:

- ***Phishing**
- Vishing
- Tailgating
- Pharming

Explanation:

Phishing aims to elicit information or get a target to perform specific actions. The threat actor attempts to deceive the user by sending a fake executable file via email and persuading them to input their network password.

Vishing is a type of social engineering attack conducted via voice channels. There is no indication of voice communication in the scenario.

Tailgating is a physical security attack where an unauthorized person accesses a restricted area by following an authorized person through a secure access point.

Pharming involves redirecting users to a fake website and hijacking the Domain Name System (DNS). There is no indication of pharming since the employee received an email with an executable file to obtain a password.

q_social_engr_tqs_pretexting_01_secp8

What social engineering technique occurs when a threat actor assumes the identity of a remote sales representative and contacts the help desk to urgently or authoritatively obtain login credentials for remote access?

Answers:

- ***Pretexting**
- Phishing
- Watering hole attack

- Brand impersonation

Explanation:

Pretexting is a social engineering tactic where a threat actor deceives a target into sharing sensitive information.

Phishing involves sending deceptive emails with disguised links to trick users into interacting with a spoofed website. However, it does not explicitly include posing as a sales representative or contacting the help desk.

A watering hole attack infiltrates a commonly used third-party website to infect devices without posing as a sales agent or contacting support.

Brand impersonation is copying a company's logos and format to create a convincing phishing message or website. It doesn't involve pretending to be a representative or contacting the help desk.

q_social_engr_tqs_pretexting_02_secp8

The cybersecurity manager of a large organization is investigating a recent security breach that occurred during office hours. Investigatory research shows that the suspect convinced the janitor to let them inside the building because they had forgotten their badge at home.

Once inside, the suspect pulled the fire alarm and accessed the building's network room amongst the chaos. The intruder then attached a monitoring device to a network port before escaping unnoticed.

Which of the following is the social engineering technique the threat actor employed in this scenario?

Answers:

- ***Pretexting**
- Vishing
- Impersonation
- Pharming

Explanation:

The threat actor used pretexting as a social engineering technique. They triggered a fire alarm to create a distraction and a convincing pretext to gain physical access to the network room, attaching a monitoring device to a network port.

Vishing is a phishing attack conducted through a voice channel (telephone or VoIP). There is no involvement of voice communication in this scenario, just a physical intrusion.

Impersonation with persuasive techniques involves convincing the target that the request is natural and without refusal. In this scenario, there is no indication of the threat actor persuading someone to act.

Pharming is a technique that redirects users from a legitimate website to a malicious one. It is irrelevant to this physical intrusion scenario.

q_social_engr_tqs_smishing_secp8

A text message purporting to be from a user's bank requests the recipient to click on a link to verify a recent transaction and provide security details.

What BEST describes the type of attack used?

Answers:

- ***SMiShing**
- Vishing
- Phishing
- Watering hole attack

Explanation:

SMiShing is a social engineering attack that uses text messages to trick people into sharing sensitive information.

Vishing is a type of phishing attack conducted through a voice channel (telephone or VoIP). The scenario explicitly describes a text message-based attack, not a voice phishing phone call.

Phishing is a social engineering technique that tricks the target into interacting with a malicious resource disguised as a trusted one, often using email as the vector. Similarly, the scenario involves a text message rather than an email.

A watering hole attack is a social engineering attack where the threat actor targets a specific group by compromising a website or online platform frequently visited by the target group.

q_social_engr_tqs_social_engineering_secp8

A threat actor poses as a remote sales representative and contacts the help desk of CloudSecure. The threat actor claims to need assistance setting up remote access.

Through a series of convincing phone calls, the threat actor obtains the name and address of the remote access server and a login credential.

What type of attack does this scenario illustrate?

Answers:

- ***Social engineering**
- Phishing
- Man-in-the-middle
- Denial-of-service

Explanation:

The stated scenario illustrates a social engineering attack where the threat actor deceives the help desk into providing sensitive information through several persuasive phone calls.

Phishing attacks involve tricking users into revealing sensitive information through email or messages, but the stated scenario focuses on several deceptive phone calls.

A man-in-the-middle attack intercepts and relays communication, which is not present in the stated scenario. The attack relies on social engineering tactics instead.

A denial-of-service attack will typically cause service disruptions. The stated scenario does not mention anything about service disruptions.

q_social_engr_tqs_typosquatting_01_secp8

An employee at a crypto security company receives an email that appears to be internal to the IT department. The email informs the employee to update the login credentials immediately to prevent account suspension.

The "From" field in the email displays "it_support@cryptosecure.com." However, upon closer inspection, the employee notices the slightly misspelled domain name as "crypt0secure.com."

What technique is the threat actor using in this phishing attempt? (Select two.)

Answers:

- ***Typosquatting**
- ***Spoofing**
- Pharming
- Brand impersonation
- Pretexting

Explanation:

The threat actor is using typosquatting by registering a domain name that is like a legitimate one, hoping users will not notice the difference and assume they are receiving emails from a known source.

The scenario also involves spoofing by impersonating a trusted source, deceiving the target. The company's spoofed email address makes it appear that the email is from the IT department.

Pharming involves redirecting users to a fake website and hijacking the Domain Name System (DNS). There is no pharming since the email has a deceptive domain name and no redirection.

Brand impersonation is committing resources to accurately duplicate a company's logos and formatting to make a phishing message or pharming website a visually compelling fake.

Pretexting is a social engineering tactic where a threat actor deceives a target into sharing sensitive information, and does not directly apply to this particular scenario.

q_social_engr_tqs_typosquatting_02_secp8

Employees at CloudCom receive a suspicious email claiming to be from "CloudCom Support," informing employees that their passwords need to be reset urgently due to a security breach. The email includes a link to a login page that looks identical to CloudCom's official site.

What type of social engineering attack does this scenario exemplify?

Answers:

- ***Typosquatting**
- Phishing
- SMiShing
- Watering hole attack

Explanation:

Typosquatting registers domains like legitimate ones, making users believe they're accessing a trusted site. The attacker creates a hijacked subdomain using the primary domain of a trusted cloud provider. Employees may fall victim to this attack if they overlook minor differences.

The attack uses a deceptive email and fraudulent login page, focusing on typosquatting. Phishing involves email or text-based attacks that trick the target into interacting with a malicious resource disguised as trusted.

While SMiShing involves phishing attacks using SMS texts, the scenario explicitly describes an email-based attack, not a text message.

A watering hole attack targets a specific group by compromising a website visited by the target group. In this scenario, there is no indication of a compromised website.

q_social_engr_tqs_vishing_01_secp8

What is the term for a phishing attack conducted through a voice channel, such as a phone call?

Answers:

- ***Vishing**
- Phishing
- SMiShing
- Pharming

Explanation:

Vishing refers to a phishing attack conducted through a voice channel, typically over the phone or VoIP, where the threat actor persuades the target to reveal sensitive information.

Phishing is a social engineering technique that tricks the target into interacting with a malicious resource disguised as a trusted one, often using email as the vector.

SMiShing stands for ""SMS Phishing,"" which involves sending deceptive text messages (SMS) to trick the recipient into interacting with a malicious resource or providing sensitive information.

Pharming is a technique that redirects users from a legitimate website to a malicious one; unrelated to voice channels.

q_social_engr_tqs_vishing_02_secp8

Which of the following social engineering attacks uses voice over IP (VoIP) to gain sensitive information?

Answers:

- ***Vishing**
- Spear phishing
- Shoulder surfing
- Hoax

Explanation:

Vishing is a social engineering attack that uses voice over IP (VoIP) to gain sensitive information. The term is a combination of voice and phishing.

Shoulder surfing involves looking over someone's shoulder while that person works on a computer or reviews documents.

In spear phishing, attackers gather information about the victim, such as identifying which online banks they use. They then send phishing emails for that specific bank.

Email hoaxes are often easy to spot because of the bad spelling and terrible grammar.

q_social_engr_tqs_vishing_03_secp8

A representative at a company reports receiving numerous unsolicited phone calls seeking banking information for a credit report.

Which social engineering variant is the finance director experiencing?

Answers:

- ***Vishing**
- Spear phishing
- Whaling
- SMiShing

Explanation:

Vishing is a type of phishing attack conducted through a voice channel (telephone or Voice over Internet Protocol, for instance) such as someone purporting to represent the bank or another official institution calling targets.

Spear phishing is a phishing scam where the attacker has some information that is more likely to fool an individual target by the attack.

Whaling is a type of spear phishing attack explicitly directed against the upper levels of management in an organization.

SMiShing is an attack that refers to using simple message service (SMS) text communications as the vector.

q_social_engr_tqs_voice_secp8

Vishing is a type of phishing attack that uses which kind of vector?

Answers:

- ***Voice or telephone**
- SMiShing
- Pharming
- Spear phishing

Explanation:

Vishing attacks are a social engineering tactic that uses the voice or a telephone network to deceive victims and gather sensitive information.

SMiShing is a type of phishing attack that uses simple message service (SMS) text communications as the vector.

Vishing attacks are voice-based, while pharming attacks manipulate Domain Name System (DNS) settings to redirect website traffic to fraudulent websites, making them unrelated in their methods and targets.

Spear phishing is not a type of vector but a highly targeted social engineering attack focusing on a specific individual within a company.

q_social_engr_tqs_watering_hole_01_secp8

Which of the following is a passive computer attack technique in which an attacker anticipates or observes the websites an organization uses often and infects them with malware?

Answers:

- ***Watering hole**
- Pretexting
- Social networking
- Typosquatting

Explanation:

A watering hole is a passive computer attack technique in which an attacker anticipates or observes the websites an organization uses often and infects them with malware. Members of the targeted group can then become infected.

Pretexting is using a fictitious scenario to persuade someone to perform an action or give information they aren't authorized to share.

Social networking involves using social media platforms to steal identities and information. Also, many attackers use social media to scam users. These scams are designed to entice the user to click a link that brings up a malicious site the attacker controls. Usually, the site requests personal information and sensitive data, such as an email address or credit card number.

Typo squatting, also called URL hijacking, relies on mistakes, such as typos made by users inputting a website address into a web browser. When a user enters an incorrect website address, the squatter may lead them to any URL.

q_social_engr_tqs_watering_hole_02_secp8

Which of the following is an example of a watering hole attack?

Answers:

- ***Targeting a group of individuals who frequent an unsecured third-party website to compromise their computers to gain access.**
- Sending deceptive emails to trick users into clicking on malicious links.
- Exploiting weak login credentials to gain unauthorized network access.
- Installing malicious software through a fake antivirus program.

Explanation:

Injecting malware into a popular website that a specific user group frequently visits is an example of a watering hole attack. The attack leverages users' trust in the site to compromise their devices.

Sending deceptive emails is a phishing attack, not a watering hole attack, where victims, lured through deceptive emails, interact with malicious content.

Exploiting login credentials is a network breach through weak credentials, not a watering hole attack, which typically involves compromising a website rather than the network.

Installing malicious software describes malware installation through a fake antivirus program, not a watering hole attack targeting specific websites, not software downloads.

q_social_engr_tqs_watering_hole_03_secp8

What social engineering attack relies on targeting individuals who frequent an unsecured third-party website to compromise their computers and gain access to a specific organization's systems?

Answers:

- ***Watering hole**
- Pharming
- Spear phishing
- Impersonation

Explanation:

A watering hole attack is a social engineering technique where the attacker identifies a popular and frequently visited website used by the target group and compromises that website with exploit code. Their computers become infected when target group members visit the website, and the attacker can then use this foothold to penetrate the organization's systems.

Pharming involves redirecting users from legitimate websites to malicious ones by corrupting the victim's computer's name resolution process. It is not specific to targeting a group of individuals.

Spear phishing is a highly targeted social engineering attack focusing on a specific individual within a company.

Impersonation involves pretending to be someone else to deceive the target.

q_social_engr_tqs_whaling_secp8

An attack that targets senior executives and high-profile victims is referred to as what?

Answers:

- ***Whaling**
- Vishing
- Pharming
- Scrubbing

Explanation:

Whaling is another form of phishing that targets senior executives and high-profile victims.

Pharming involves the attacker executing malicious programs on the target's computer so that when the user enters any URL, it redirects traffic to the attacker's malicious website.

Vishing is like phishing, but instead of an email, the attacker uses voice over IP (VoIP) to gain sensitive information. The term is a combination of voice and phishing.

Scrubbing is one of the most common ways to pick a lock.

2.3 Malware

As you study this section, answer the following questions:

- What is the difference between a virus and a worm?
- Which types of malware typically use email to spread?
- What does it mean for software to be quarantined?
- Why is it a good practice to show file extensions?
- What must you do to ensure that you are protected from the latest virus variations?

In this section, you will learn to:

- Implement malware protections.
- Use Windows security.
- Configure Windows Defender protections to secure a network from malware.

Key terms for this section include the following:

Term	Definition
Malware	Software that serves a malicious purpose, typically installed without the user's consent (or knowledge).
Trojan	A malicious software program hidden within an innocuous-seeming piece of software. Usually, the Trojan is used to try to compromise the security of the target computer.
Potentially unwanted programs (PUPs)/potentially unwanted applications (PUAs)	Software that cannot definitively be classed as malicious, but may not have been chosen or wanted by the user.
Virus	Malicious code inserted into an executable file image. The malicious code is executed when the file is run and can deliver a payload, such as attempting to infect other files.
Malicious process	A process executed without proper authorization from the system owner for the purpose of damaging or compromising the system.
Worm	A type of malware that replicates between processes in system memory and can spread over client/server network connections.
Shellcode	A lightweight block of malicious code that exploits a software vulnerability to gain initial access to a victim system.

Advanced persistent threat (APT)	An attacker's ability to obtain, maintain, and diversify access to network systems using exploits and malware.
Adware	Software that records information about a PC and its user. Adware is used to describe software that the user has acknowledged can record information about their habits.
Spyware	Software that records information about a PC and its users, often installed without the user's consent.
Keylogger	Malicious software or hardware that can record user keystrokes.
Backdoor	A mechanism for gaining access to a computer that bypasses or subverts the normal method of authentication.
Remote access Trojan (RAT)	Malware that creates a backdoor remote administration channel to allow a threat actor to access and control the infected host.
Botnet	A group of hosts or devices that has been infected by a control program called a bot, which enables attackers to exploit the hosts to mount attacks.
Command and control (C2 or C&C)	Infrastructure of hosts and services with which attackers direct, distribute, and control malware over botnets.
Covert channel	A type of attack that subverts network security systems and policies to transfer data without authorization or detection.
Internet Relay Chat (IRC)	A group communications protocol that enables users to chat, send private messages, and share files.
Rootkit	Class of malware that modifies system files, often at the kernel level, to conceal its presence.
Ransomware	Malware that tries to extort money from the victim by blocking normal operation of a computer and/or encrypting the victim’s files and demanding payment.
Crypto-mining	Malware that hijacks computer resources to create cryptocurrency.

Logic bomb	A malicious program or script that is set to run under particular circumstances or in response to a defined event.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.4 Given a scenario, analyze indicators of malicious activity.
	• Malware attacks
	  ○ Ransomware
	  ○ Trojan
	  ○ Worm
	  ○ Spyware
	  ○ Bloatware
	  ○ Virus
	  ○ Keylogger
	  ○ Logic bomb
	  ○ Rootkit

TestOut Security Pro	3.1 Harden computer systems 3.1.2 Configure anti-virus protection

2.3.1 Malware (Lesson Video)

Transcript:

Malware is software designed to infiltrate or damage a computer system without the owner's consent or knowledge. It often attempts to hide itself. And some malware even takes control of the computer system.

There's a long list of malware types. I'm going to tell you about just a few of them.
First, let me introduce you to worms. A worm propagates without a file. It travels across computer networks without any user assistance, and it automatically replicates itself. Worms are one of the deadliest malware types because they can quickly spread to millions of computers. They usually take advantage of unpatched vulnerabilities in computer systems. To avoid a worm infection, keep your systems patched.
Next, there's the Trojan. The Trojan virus is named after the Trojan horse in Greek mythology. It seems benign, but it's not. It usually comes in the form of a useful application hiding malicious code. It may be explicit or run invisibly. Trojans don't self-replicate like worms, and they aren't attached to a file like a virus. But they can create back doors for an attacker, allowing remote control of the infected machine.
The Remote Access Trojan, or RAT, is a common virus. A RAT is designed to give the attacker remote desktop access, including a GUI. It gives the hacker complete control over the system and allows them to view all communications, webcam footage, files, and emails.
Since a RAT requires the hacker to connect to the user through a port that isn't commonly used for internet access, if the user is behind a firewall, this often prevents the hacker from connecting. If the hacker can get around the firewall and open that port, then they're able connect.
Next, there are rootkits. This malicious software hides in the system and alters the system's processes and registry entries. A rootkit is almost invisible. It resides below antivirus software detection. This makes it difficult to detect and very dangerous. Special Anti-rootkit software can help detect rootkits.

Let's move on to the logic bomb. This is a piece of malicious code designed to execute only under pre-defined conditions. It'll lay dormant until the pre-defined condition is met, such as a certain time or date. Logic bombs can range from benign, such as pranks, to dangerous, such as driver formatters. One example is a piece of code that lies dormant until it executes once per year, perhaps on someone's birthday. Another example is a user opening a certain program that activates it, deleting the contents of disk drives.
Next is spyware. This is computer software installed without the user's consent or knowledge. It's designed to intercept data or take partial control of the user's computer. Spyware often collects personal information about the user, like internet surfing habits and passwords. It uses tracking cookies to collect information about the user's activities and report it to the hacker. It can also interfere with a user's control in other ways, such as installing additional software, changing computer settings, and redirecting web browsers.
An important malware to know is ransomware. It allows a hacker to gain access to a system, plant a virus that encrypts all user data, and then demand a payment for decrypting the data. If the victim doesn't pay, the hacker threatens to destroy the data. But there's no guarantee that the hacker will actually send the decryption key once the victim's paid for it. There's also no guarantee that the victim will pay the ransom. If they have a backup still intact, they may choose not to pay.
Because of this, ransomware is becoming less prevalent and cryptomalware is gaining popularity. The difference between ransomware and cryptomalware is that crypto operates quietly in the background and keeps going indefinitely until it's noticed. This is called cryptomining. Mining cryptocurrency wears down a system substantially. It eats up bandwidth and processing power, which slows the system down and affects productivity. If it goes on long enough it can even cause graphic cards to die, processors to burn out, and memory to act unpredictably.
Another interesting malware is the fileless virus. It uses legitimate programs to infect a computer. Because it doesn't rely on files, it leaves no footprint, making it undetectable by most antivirus, whitelisting, and other traditional endpoint security solutions. Fileless malware isn't a traditional virus, but it works in a similar way by operating in memory. It never even touches the hard drive. Many hackers use social engineering to get users to click on a link in a phishing email. When the web page opens, the virus gets into the inner recesses of a trusted application such as PowerShell or Windows script host executables.
Now let's talk about how you can prevent falling victim to malicious software. First, use the latest version and patch level for your web browser. Then install the latest patches for your operating system. You should also install antivirus, anti-spyware, anti-rootkit, and personal firewall software. On all anti-malware software, be sure to keep your definitions up to date for the latest threats. Use a pop-up blocker to prevent adware. Use software to control cookies on your system. Schedule regular scans to constantly detect malware.
If a scan does detect malware, quarantine or delete the malicious software. Malware may cause permanent damage to the system. Even though anti-malware software can scan, detect, and delete malware for you, the system may still need repairs. You might have to reinstall applications or even reinstall the entire operating system.
Many companies image their operating systems, so if malware is detected, they'll simply send down a brand new image that overrides the entire operating system.
All right, that's it for this lesson. In this video, we talked about different types of malware and how they work. We also learned how to prevent and repair malware in case your system or network becomes infected.

2.3.2 Malware Facts

By classifying the various types of malware and identifying the signs of infection, security teams are better prepared to remediate compromised systems or prevent malware from executing in the first place.

This lesson covers the following topics:

- Malware classification
- Computer viruses
- Spyware and Keyloggers
- Backdoors and Remote Access Trojans
- Rootkits
- Ransomware, Crypto-Malware, and Logic Bombs
- Crypto-Ransomware
- Logic Bombs

Malware Classification

Many of the intrusion attempts perpetrated against computer networks depend on the use of malicious software or malware. Malware is simply defined as software that does something bad from the perspective of the system owner. A complicating factor with malware classification is the degree to which its installation is expected or tolerated by the user.

Some malware classifications, such as Trojan, virus, and worm, focus on the vector used by the malware. The vector is the method by which the malware executes on a computer and potentially spreads to other network hosts. The following categories describe some types of malware according to vector:

- **Viruses and worms** represent some of the first types of malware and spread without any authorization from the user by being concealed within the executable code of another process. These processes are described as being infected with malware.
- **Trojan** refers to malware concealed within an installer package for software that appears to be legitimate. This type of malware does not seek any type of consent for installation and is actively designed to operate secretly.
- **Potentially unwanted programs (PUPs)/Potentially unwanted applications (PUAs)** are software installed alongside a package selected by the user or perhaps bundled with a new computer system. Unlike a Trojan, the presence of a PUP is not automatically regarded as malicious. It may have been installed without active consent or with consent from a purposefully confusing license agreement. This type of software is sometimes described as *grayware* rather than malware. It can also be referred to as bloatware.

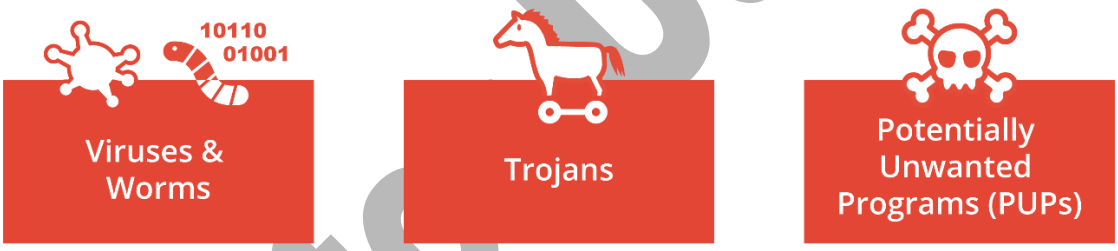

Malware classification by vector.

Other classifications are based on the payload delivered by the malware. The *payload* is an action performed by the malware other than simply replicating or persisting on a host. Examples of payload classifications include spyware, rootkit, remote access Trojan (RAT), and ransomware.

Computer Viruses

A computer virus is a type of malware designed to replicate and spread from computer to computer, usually by "infecting" executable applications or program code. There are several different types of viruses, and they are generally classified by the different types of file or media that they infect:

- **Non-resident/file infector** — the virus is contained within a host executable file and runs with the host process. The virus will try to infect other process images on persistent storage and perform other payload actions. It then passes control back to the host program.
- **Memory resident** — when the host file is executed, the virus creates a new process for itself in memory. The malicious process remains in memory, even if the host process is terminated.
- **Boot** — the virus code is written to the disk boot sector or the partition table of a fixed disk or USB media and executes as a memory-resident process when the OS starts, or the media is attached to the computer.
- **Script and macro viruses** — the malware uses the programming features available in local scripting engines for the OS and/or browser, such as PowerShell, Windows Management Instrumentation (WMI),

JavaScript, Microsoft Office documents with Visual Basic for Applications (VBA) code enabled, or PDF documents with JavaScript enabled.

In addition, the term "multipartite" is used for viruses that use multiple vectors, and the term "polymorphic" is used for viruses that can dynamically change or obfuscate their code to evade detection.

What these types of viruses have in common is that they must infect a host file or media. An infected file can be distributed through any normal means—on a disk, on a network, as an attachment to an email or social media post, or as a download from a website.

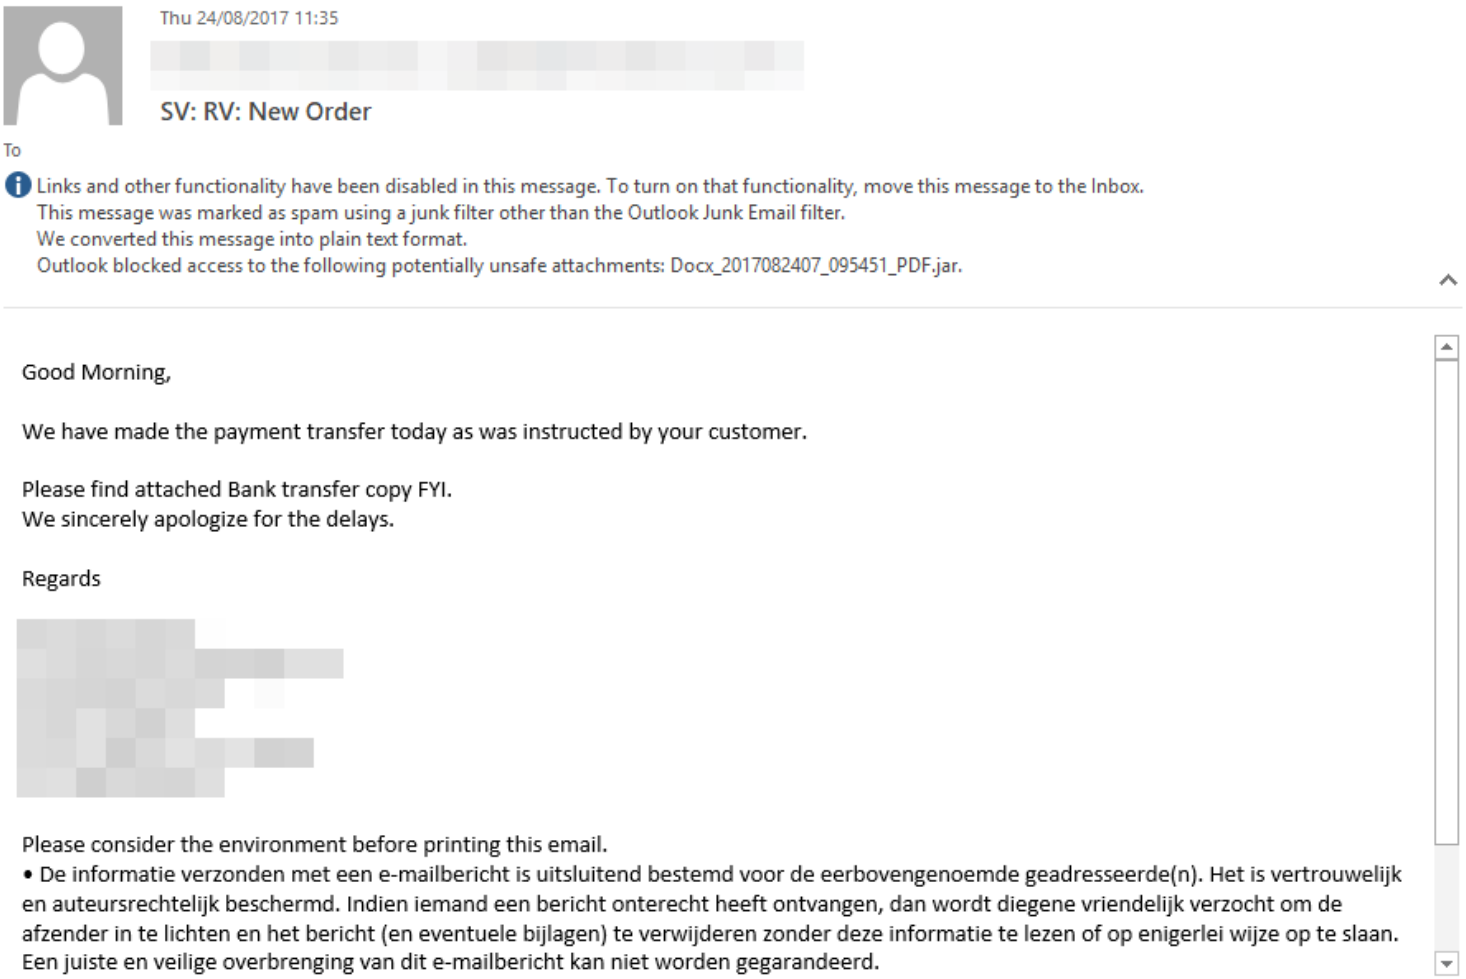

Unsafe attachment detected by Outlook's mail filter—The "double" file extension is an unsophisticated attempt to fool any user not already alerted by the use of both English and German in the message text. (Screenshot used with permission from Microsoft.)

A computer worm is memory-resident malware that can run without user intervention and replicate over network resources. A virus is executed only when the user performs an action such as downloading and running an infected executable process, attaching an infected USB stick, or opening an infected document with macros or scripting enabled. By contrast, a worm can execute by exploiting a vulnerability in a process when the user browses a website, runs a vulnerable server application, or is connected to an infected file share. For example, the Code Red worm was able to infect early versions of Microsoft's IIS web server software via a buffer overflow vulnerability. It then scanned randomly generated IP ranges to try and infect other vulnerable IIS servers.

The primary effect of the first type of computer worm is to rapidly consume network bandwidth as the worm replicates. A worm may also be able to crash an operating system or server application, performing a denial of service attack. Also, like viruses, worms can carry a payload that can be written to perform any type of malicious action.

The Conficker worm illustrated the potential for remote code execution and memory-resident malware to effect highly potent attacks. As malware has continued to be developed for criminal intent and security software became better able to detect and block static threats, malware code and techniques have become more sophisticated. The term "fileless" has gained prominence to refer to these modern types of malware. Fileless is not a definitive classification, but it describes a collection of common behaviors and techniques:

- Fileless malware does not write its code to disk. The malware uses memory-resident techniques to run in its own process, within a host process or dynamic link library (DLL), or a scripting host. This does not mean that there is no disk activity at all, however. The malware may change registry values to achieve persistence (executing if the host computer is restarted). The initial execution of the malware may also depend on the user running a downloaded script, file attachment, or Trojan software package.
- Fileless malware uses lightweight shellcode to achieve a backdoor mechanism on the host. The shellcode is easy to recompile in an obfuscated form to evade detection by scanners. It is then able to download additional packages or payloads to achieve the threat actor's objectives. These packages can also be obfuscated, streamed, and compiled on the fly to evade automated detection.
- Fileless malware may use "live off the land" techniques rather than compiled executables to evade detection. This means that the malware code uses legitimate system scripting tools, notably PowerShell and Windows Management Instrumentation (WMI), to execute payload actions. If they can be executed with sufficient permissions, these environments provide all the tools the attacker needs to perform scanning, reconfigure settings, and exfiltrate data.

The terms " advanced persistent threat (APT) " and "advanced volatile threat (AVT)" can be used to describe this general class of modern fileless/live-off-the-land malware. Another useful classification is a low-observable characteristic (LOC) attack. The exact classification is less important than the realization that adversaries can use any variety of coding tricks to effect intrusions and that their tactics, techniques, and procedures to evade detection are continually evolving.

Spyware and Keyloggers

The first viruses and worms focused on the destructive potential of being able to replicate. As the profitable uses of this software became apparent, however, they started to be coded with payloads designed to facilitate intrusion, fraud, and data theft. Bloatware and malware can be used for different levels of monitoring:

- **Tracking cookies** — a cookie is a plaintext file, not malware, but if permitted by browser settings, third-party cookies can be used to record web activity, track the user's IP address, and harvest various other metadata, such as search queries and information about the browser software and configuration. Tracking cookies are created by adverts and analytics widgets embedded into many websites.
- **Supercookies and beacons** — as browser software gives the user some control over what cookies to accept, web marketing companies have come up with alternative ways to implement tracking that are difficult to disable. A supercookie is a means of storing tracking data in a non-regular way, such as saving it to a cache without declaring the data to be a cookie or encoding data into header requests. A beacon is a single-pixel image embedded into a website. While invisible to the user, the browser must request to download the pixel to load the site, giving the beacon host the opportunity to collect metadata, perform browser fingerprinting, and potentially run tracking scripts.
- **Adware** — this is a class of PUP/bloatware that performs browser reconfigurations, such as allowing tracking cookies, changing default search providers, opening sponsor's pages at startup, adding bookmarks, and so on. Adware may be installed as a program or browser extension/plug-in.
- **Spyware** — this is malware that can perform adware-like tracking but also monitor local application activity, take screenshots, and activate recording devices, such as a microphone or webcam. Another spyware technique is to perform DNS redirection to pharming sites.
- **Keylogger** — this is spyware that actively attempts to steal confidential information by recording keystrokes. The attacker will usually hope to discover passwords or credit card data.

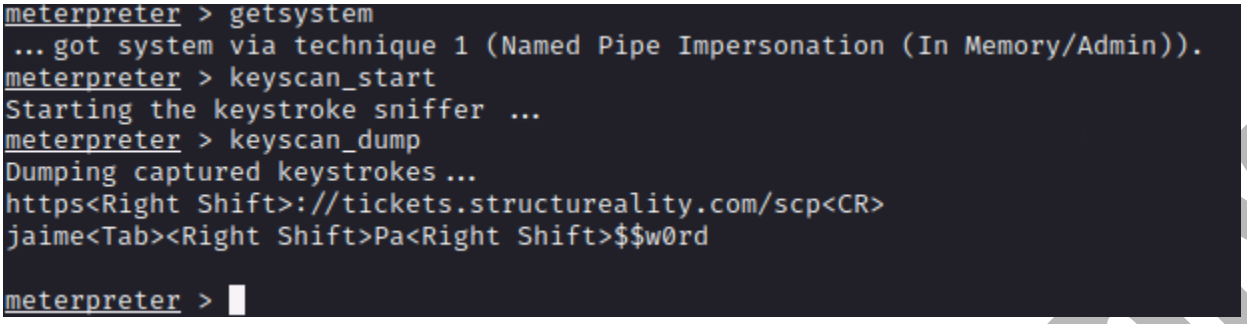

Using the Metasploit Meterpreter remote access tool to dump keystrokes from the victim machine, revealing the password used to access a web app.

Keyloggers are not only implemented as software. A malicious script can transmit key presses to a third-party website. There are also hardware devices to capture key presses to a modified USB adapter inserted between the keyboard and the port. Such devices can store data locally or come with Wi-Fi connectivity to send data to a covert access point. Other attacks include wireless sniffers to record key press data, overlay ATM pin pads, and so on.

Backdoors and Remote Access Trojans

Any type of access method to a host that circumvents the usual authentication method and gives the remote user administrative control can be referred to as a backdoor . A remote access Trojan (RAT) is backdoor malware that mimics the functionality of legitimate remote control programs but is designed specifically to operate covertly. Once the RAT is installed, it allows the threat actor to access the host, upload files, and install software or use "live off the land" techniques to effect further compromises. In this context, RAT can also stand for remote administration tool. A host that is under malicious control is sometimes described as a zombie.

A compromised host can be installed with one or more bots. A bot is an automated script or tool that performs some malicious activity. A group of bots that are all under the control of the same malware instance can be manipulated as a botnet by the herder program. A botnet can be used for many types of malicious purposes, including triggering distributed denial of service (DDoS) attacks, launching spam campaigns, or performing cryptomining.

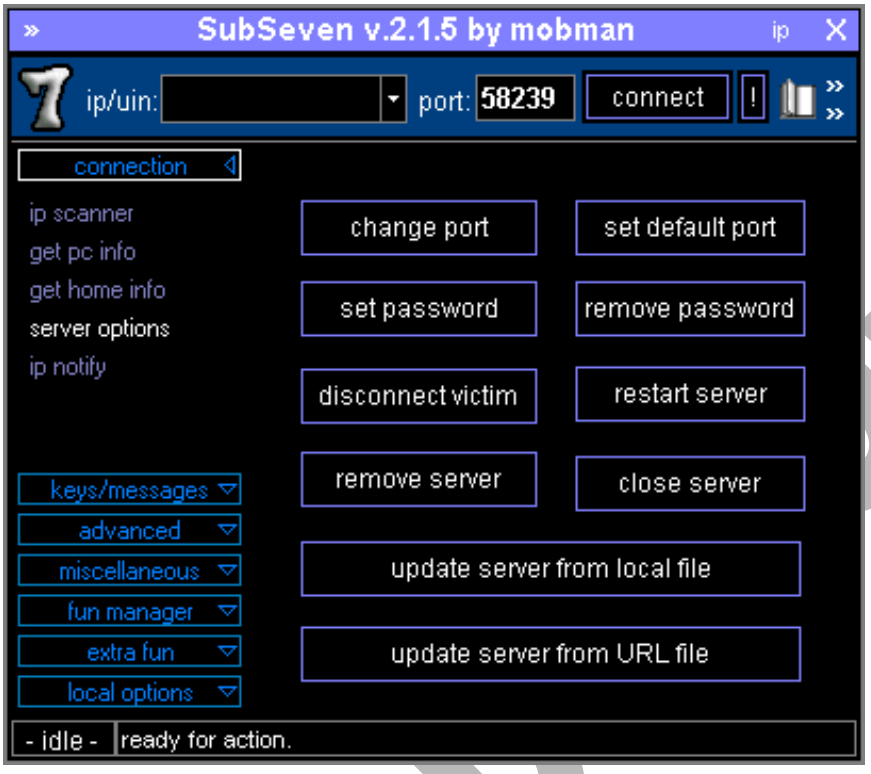

SubSeven RAT. (Screenshot used with permission from Wikimedia Commons by CCAS4.0 International.)

Whether a backdoor is used as a standalone intrusion mechanism or to manage bots, the threat actor must establish a connection from the compromised host to a command and control (C2 or C&C) host or network. This network connection is usually the best way to identify the presence of a RAT, backdoor, or bot. There are many means of implementing a C&C network as a covert channel to evade detection and filtering. Historically, the Internet Relay Chat (IRC) protocol was popular. Modern methods are more likely to use command sequences embedded in HTTPS or DNS traffic.

Backdoors can be created in other ways than by infection by malware. Programmers may create backdoors in software applications for testing and development that are subsequently not removed when the application is deployed. Backdoors are also created by misconfiguration of software or hardware that allows access to unauthorized users.

Rootkits

In Windows, Trojan malware that depends on manual execution by the logged-on user inherits the privileges of that user account. If the account has only standard permissions, the malware can only add, change, or delete files in the user's profile and run only apps and commands that the user is permitted to.

If the malware attempts to change system-wide files or settings, it requires local administrator-level privileges. To obtain those through manual installation or execution, the user must be confident enough in the Trojan package to confirm the User Account Control (UAC) prompt or enter the credentials for an administrative user.

If the malware gains local administrator-level privileges, there are still protections in Windows to mitigate abuse of these permissions. Critical processes run with a higher level of privilege called SYSTEM. Consequently, Trojans installed or executed with local administrator privileges cannot conceal their presence entirely and will show up as a running process or service. Often, the process image name is configured to resemble a genuine executable or library to avoid detection. For example, a Trojan may use the filename "rund1132.exe" to masquerade as "rundll32.exe." To ensure persistence (running when the computer is restarted), the Trojan may have to use a registry entry or create itself as a service, which can usually be detected easily.

If the malware can be delivered as the payload for an exploit of a severe vulnerability, it may be able to execute without requiring any authorization using SYSTEM privileges. Alternatively, the malware may be able to use an exploit to escalate privileges to SYSTEM level after installation. Malware running with this level of privilege is referred to as a rootkit . The term derives from UNIX/Linux, where any process running as the root superuser account has unrestricted access to everything from the root of the file system down.

In theory, there is nothing about the system that a rootkit could not change. In practice, Windows uses other mechanisms to prevent misuse of kernel processes, such as code signing. Consequently, what a rootkit can do depends largely on adversary capability and level of effort. When dealing with a rootkit, you should be aware that there is the possibility that it can compromise system files and programming interfaces, so that local shell processes, such as Explorer, taskmgr, or tasklist on Windows or ps or top on Linux, plus port scanning tools, such as netstat, no longer reveal its presence (at least, if run from the infected machine). A rootkit may also contain tools for cleaning system logs, further concealing its presence.

Software processes can run in one of several "rings." Ring 0 is the most privileged (it provides direct access to hardware) and should be reserved for kernel processes only. Ring 3 is where user-mode processes run; drivers and I/O processes may run in Ring 1 or Ring 2. This architecture can also be complicated by the use of virtualization.

There are also examples of rootkits that can reside in firmware (either the computer firmware or the firmware of any sort of adapter card, hard drive, removable drive, or peripheral device). These can survive any attempt to remove the rootkit by formatting the drive and reinstalling the OS. For example, the US intelligence agencies have developed DarkMatter and Quark Matter UEFI rootkits targeting the firmware on Apple Macbook laptops.

Ransomware, Crypto-Malware, and Logic Bombs

Ransomware is a type of malware that tries to extort money from the victim by making the victim’s computer and/or data files unavailable and demanding payment. One class of ransomware will display threatening messages, such as requiring Windows to be reactivated or suggesting that the computer has been locked by the police because it was used to view child pornography or for terrorism. This may block access to the file system by installing a different shell program, but this sort of attack is usually relatively simple to fix.

Ransomware uses payment methods, such as wire transfer, cryptocurrency, or premium rate phone lines, to allow the attacker to extort money without revealing their identity or being traced by local law enforcement.

WannaCry ransomware. (Image by Wikimedia Commons.)

Scareware refers to malware that displays alarming messages, often disguised to look like genuine OS alert boxes. Scareware attempts to alarm the user by suggesting that the computer is infected or has been hijacked.

Crypto-Ransomware

The crypto class of ransomware attempts to encrypt data files on any fixed, removable, and network drives. If the attack is successful, the user will be unable to access the files without obtaining the private encryption key, which is held by the attacker. If successful, this sort of attack is extremely difficult to mitigate unless the user has backups of the encrypted files. One example of crypto-ransomware is CryptoLocker, a Trojan that searches for files to encrypt and then prompts the victim to pay a sum of money before a certain countdown time, after which the malware destroys the key that allows the decryption.

Cryptojacking Malware

Another type of crypto-malware hijacks the resources of the host to perform cryptocurrency mining. This is referred to as crypto-mining , and malware that performs crypto-mining maliciously is classed as *cryptojacking* . The total number of coins within a cryptocurrency is limited by the difficulty of performing the calculations necessary to mint a new digital coin. Consequently, new coins can be very valuable, but it takes enormous computing resources to discover them. Cryptojacking is often performed across botnets.

Logic Bombs

Some types of malware do not trigger automatically. Having infected a system, they wait for a preconfigured time or date (time bomb) or a system or user event (logic bomb). Logic bombs also need not be malware code. A typical example is a disgruntled systems administrator who leaves a scripted trap that runs in the event their account is deleted or disabled. Antivirus software is unlikely to detect this kind of malicious script or program. This type of trap is also referred to as a *mine* .

2.3.3 Malware Protection Facts

Malware is software designed to infiltrate or damage a computer system without the owner's consent or knowledge. Some malware attempts to hide itself or can provide a remote attacker control of the computer system.

This lesson covers the following topics:

- Malware prevention
- Malware recovery

Malware Prevention

Regardless of the type of malware, there are some common things you can do to prevent malware infection:

Option	Description
Patches	As corporations find errors, they regularly release them in updates and patches. Ensure that you are at the current version of the following: • Hardware Firmware Updates • Server and Client Operating Systems • Applications, especially your web browser • Malware protection software
Anti-malware	Install antivirus or anti-malware software to protect against known infections. Consider the following: • Choose anti-malware software from a reputable company. Do not let scareware fool you into purchasing a product that may not work or is actually a virus. • Keep definition files up-to-date. • Enable antivirus scanning for email attachments. • Enable scanning of files downloaded from the internet. • Enable antivirus scanning for all removable storage, such as USB flash drives and CD-ROMs. • Schedule regular scans.
Browser Settings	Utilize browser settings to help prevent infection, such as: • Enable a pop-up blocker. • Enable privacy controls. • Monitor and manage cookies. • Regularly clear the cache/history.

Option	Description
	• Disable or configure Autocomplete settings not to store sensitive information.
Firewall	Modern operating systems include a local firewall. This will help protect against unauthorized connections to the local host and attempts by malicious software to make outgoing connections to command and control servers.
Training	Educate end users about common safety tips, such as: • Examine and email attachments for incorrect file extensions and confirm with the sender what they sent if possible. • If pop-ups occur for unknown software, ask to install a new antivirus or more; do not accept and exit the dialog. Then contact IT. • Do not use social media at work where malicious links could cause the work computer to be infected. • Hover over all links to verify that the target URL is valid. • Do not use any devices that they find lying around the office. • Only use reputable software approved by the organization.
Web Filter	If there is a need to have additional control, consider installing a web filter. This provides the ability to block access to malicious sites that users might be visiting regularly.

Malware Recovery

Malware can permanently damage your system. Recovery from malware can include the following steps:

- **Remediate** problems as prompted. *Remediation* is the process of correcting problems. Most antivirus software remediates problems automatically or semi-automatically by prompting you to identify the action to take. Possible actions in response to problems are:
    - **Repair the infection** . Repair is possible for true viruses that have attached themselves to valid files. During the repair, the virus is removed, and the file is placed back in its original state when possible.
    - **Quarantine the file** . Quarantine moves the infected file to a secure folder where it cannot open or run normally. You might quarantine an infected file that cannot be repaired to see if another tool or utility can recover the file at another time.
    - **Delete the file** . You should delete malicious files such as worms, Trojan horse programs, spyware, or adware programs.
    - Periodically review the quarantine folder and delete any files you do not want to recover.
- **Reinstall** applications or features.
- **Re-image** a machine if your organization uses imaging solutions. Re-imaging or installing from scratch is often faster and more effective than malware removal and cleanup. If an image is not available, reinstall the operating system.

2.3.4 Implementing Malware Protections (Demo Video)

Transcript:

In this demonstration, we're going look at the Windows Security app. In previous versions of Windows, this was known as Windows Defender or Windows Defender Security Center. Either way, its functionality as Windows' default anti-

malware software still remains the same. The name of the anti-malware software that's controlled by this app has also been renamed. It's called Microsoft Defender now.

Let's delve into the Windows Security app a little more. Let's come down here, type Security, and click the Windows Security app to open it.

When you start the app up, you'll see Security at a glance. They've divided the whole thing into smaller subsections for easier navigation, like Virus & threat protection, Account protection, Firewall & network protection, and so on. If you look at the bottom-right corner of each category icon, you'll see a green checkmark that indicates that each section is currently providing protection.

If, on the other hand, you see a yellow exclamation point triangle on the icon, that indicates that something might be amiss. And if you see a red X here, it means that something needs your immediate attention.

Okay, let's take a closer look at Virus & threat protection. As I mentioned before, the green checkmark tells us that no action is needed at this moment. But let's click it to take a closer look anyway.

Virus & threat protection is what Windows Defender used to be—it takes a thorough look at your computer system and checks for things that might be concerning. Here, we can look at Current threats, Protection settings, Protection updates, and Ransomware protection.

Under Current threats, we see the last time the system was scanned, how many threats were found, and how many files were scanned as well. By default, Microsoft Defender scans your system on a regular basis. But if you're concerned that some type of malware may have gotten into your system and you don't want to wait until the next scheduled scan, you can manually scan the system whenever you like by clicking Quick scan. You can change Scan options as well.

Notice that you have four options for running a manual scan. The first one is the Quick scan I mentioned, which checks your system in places where threats are commonly found. A Full scan checks every file on the system. The disadvantage of full scans, obviously, is that they take a lot longer. The third option would be to run a Custom scan to check specific files and folders. The last option is a Microsoft Defender Offline scan. This one takes about 15 minutes and restarts the system as part of the routine.

If you run a scan and do find malware, it means that your system is already infected and needs to be remediated. Obviously, it's much better to try to detect malware before it infects your system, say, as you're downloading a file from a malicious website. That's where Real-time protection comes in.

Instead of just scanning for malicious files on your hard disk drives, Real-time protection watches what you're doing and tries to detect malware as it enters the system. To turn on Real-time protection, let's come down here, to Virus & threat protection settings. Notice that it tells us that Real-time protection is currently turned on.

Understand that with this software, there are a few different ways to utilize the protection options. One is to generically scan files for characteristics that might indicate that that file is either malware or contains malware of some sort. There's also Cloud-delivered protection, which sends information to Microsoft that allows them to update their definition files and get them out faster. We'll talk about definition files shortly. You also have the option to send those malware samples to Microsoft directly. You can also use Tamper Protection, which monitors important security features and prevents them from being changed. Controlled folder access allows further protection for dedicated locations and files on your system.

Occasionally, you may need to implement an exclusion. This is any item, file, folder, or process that you want to exclude from being scanned. Take caution when you're adding exclusions, as some excluded items could contain threats that could make your system vulnerable. To add an exclusion, click the Add or remove exclusions link, click the Add an exclusion button, select the type of item to be added, browse to it, and select it.

Now let's go back and look at Virus & threat protection updates. First, it's important to know that Windows Security uses files called definitions to help detect threats and protect your system from them. Windows Defender tries to automatically download the latest definitions for you.

Under Virus & threat protection updates, it tells us the last time that the system was scanned using definitions. If we click Protection updates, we can see even more information, such as the version, when it was created, and when it was last updated. If we want to manually force an update, we can click Check for updates.

Windows Security works reasonably well. However, some folks still prefer to disable it and install a third-party anti-malware package. If you choose this route, remember that before you install the package on your system, you need to make sure that Windows Security is turned off by going to Virus & threat protection and turning Real-time protection off, as we did earlier in this demo.

This is important because if you have two anti-malware packages running on your system at the same time, they're likely to conflict. Each will probably think that the other one is a virus.

There are some good free anti-malware packages, but keep in mind that you're likely to get better service from a paid package than a free one. But at the end of the day, the important thing is that you protect your computer from malware in some way.

That's it for this demonstration on malware. We talked about implementing anti-malware software by looking at the Windows Security app, which comes with Windows by default. Then we took a closer look at the app's features in Virus & threat protection. We also talked briefly about third-party anti-malware packages.

2.3.5 Use Windows Security (Demo Video)

Transcript:

In this demonstration, we're going to practice working with Windows Defender Firewall, which is a host-based firewall, so it's implemented as software and designed to protect an individual system. It's good to know that previously, Windows Defender Firewall was referred to as just Windows Firewall. So if you're on an older operating system, it'll go by that name, but they work basically the same. Like any firewall, it's job is to prevent somebody out on the internet or on the network from initiating an unwanted connection with your system. It just shuts any threats down before they can even get in.

Okay, we can manage these settings by searching for Control Panel, selecting System and Security, and then selecting Windows Defender Firewall. You'll notice that the firewall state is currently off. This isn't good. We want to turn the firewall back on, so let's go over here and click on Turn Windows Defender Firewall on or off.
You'll see that there are two different network profiles where the firewall is applied—our private and public network settings. We can have the private settings be a little more lax because we're assuming that the company already has a firewall in place that's blocking a lot of connections. We mostly just want to protect this individual host, which is why we'd customize the public settings more granularly. This is a bigger issue because public networks—like those at restaurants, hotels, or airports—are like the wild, wild west of information technology. You have no idea who else is on there or what their intentions are.
For both of the profiles, you can see that firewall protection is currently turned off. Again, that's not a good thing. Let's change that by selecting On for both. These default settings here work fairly well for a private network, and we'll be notified if Windows Defender Firewall blocks an app that's trying to communicate through the network. Then we can just make a list of apps with permission to pass through if we need. That's appropriate to do because there are obviously certain applications we do want to be able to receive data from. However, this might not be the case with a public network. If I'm on a public network, I most likely want to turn this option on and block all incoming connections, including those on the list of allowed apps. There's probably not anyone out there on a public network that should be initiating a connection with my computer, even with an allowed app. We'll go ahead and leave this option on for public network profiles. Let's click OK, and now the Windows Defender Firewall state is set to On for incoming connections.
Let's go deeper into how we make our list of allowed apps. There may be times when you install an app, and it needs to communicate with other hosts on the network. Like you saw, we just blocked everything by default. To change that, we come over here to Allow an app or feature through Windows Defender Firewall and select that option.
Here you can see a list of applications and features that the Windows system is aware of at this point in time and whether or not they're allowed through the firewall. For example, I'll scroll down to Remote Assistance. Remote Assistance traffic is currently allowed through the host firewall for both private and public profiles. I've enabled Remote Desktop specifically because I do a lot of remoting into my machines for testing purposes. It's usually not too difficult to find an application that you want to allow to pass through. Sometimes, though, you can't find the app you want in the list because Windows hasn't recognized it yet.
If that's the case, you can come down here and click Allow another app. Now you need to locate the app yourself. I have an app here in my Downloads folder called VNC Viewer, which allows me to connect to desktops remotely not only on Windows, but also on Linux and UNIX systems that aren't designed for use with Remote Desktop. I'll go ahead and select my VNC Viewer app and click Add. It automatically selected Public, and I can manually choose Private as well. That's it! I just click OK to finalize everything. Now that we've configured the firewall settings we want, we're ready to go.
That's all for now. In this demo, we talked about how to manage Windows Defender Firewall. First, we went over how to turn the firewall on and off, and we talked about increasing the firewall's security level on different network profiles. Then we ended this lesson by going over how to add exceptions so that specific applications can pass through.

2.3.6 Configure Microsoft Defender (Simulation)

Scenario

You recognize that the threat of malware is increasing. As such, you would like to use Windows Virus & Threat Protection to protect your computer from malware.

In this lab, your task is to:

- Enable and configure Windows Virus & Threat Protection as follows:
  - Add a file exclusion for **D:\Graphics\cat.jpg** .
  - Add a process exclusion for **welcome.scr** .
  - Locate the current threat definition version number.
    - Answer Question 1.
  - Check for updates.
    - Answer Question 2.
  - Perform a quick scan.

Explanation

Complete this lab as follows:

1. Access the *Virus & threat protection* options.
   a. Right-click **Start** , and then select **Settings** .
   b. Select **Update & Security** .
   c. From the left pane, select **Windows Security** .
   d. From the right, select **Virus & threat protection** .
2. Add a file exclusion for D:\Graphics\cat.jpg.
   a. Under *Virus & threat protection settings* , select **Manage settings** .
   b. Scroll down to Exclusions and then select **Add or remove exclusions** .
   c. Select **Add an exclusion** , and then select **File** .
   d. From the left pane, browse to and select **Data (D:)** > **Graphics** .
   e. From the right pane, select **cat.jpg** , and then select **Open** .
3. Add a process exclusion for welcome.scr.
   a. From the Exclusions dialog, select **Add an exclusion** , and then select **Process** .
   b. In the *Enter process name* field, type **welcome.scr** and select **Add** .
4. Check for protection updates.
   a. From the top left of Windows Security, select the **back arrow** twice to return to the *Virus & threat protection* page.
   b. Scroll down to *Virus & threat protection updates* and then select **Check for updates** to access the Protection updates page.
   c. From the top right, select **Answer Questions** .
   d. Answer Question 1.
   e. Select **Check for updates** .
   f. Answer Question 2.
5. Perform a quick virus scan.
   a. From the top left of the Windows Security dialog, select the **back arrow** to return to the *Virus & threat protection* page.
   b. Select **Quick scan** .
   c. Wait for the scan to complete.
6. From the Lab Questions dialog , select **Score Lab** .

2.3.7 Analyze Indicators of Malware-Based Attacks

2.3.8 Practice Questions (Section Quiz)

q_mal_bloatware_secp8

Which type of malware is software installed alongside a package selected by the user or bundled with a new computer system?

Answers:

- Logic bomb
- Trojan horse
- Spyware
- ***Bloatware**

Explanation:

Bloatware (or PUP) is software installed alongside a package selected by the user or perhaps bundled with a new computer system. Unlike a Trojan, the presence of a PUP is not automatically regarded as malicious. It may have been installed without active consent or with consent from a purposefully confusing license agreement.

A logic bomb is malware that lies dormant until triggered.

A Trojan horse is a malicious program that is disguised as legitimate software.

Spyware monitors the actions performed on a machine and then sends the information back to its originating source.

q_mal_crypto-ransomware_secp8

What is the primary function of crypto-ransomware?

Answers:

- To steal sensitive information from the infected system.
- ***To encrypt files on the infected system and demand a ransom for the decryption key.**
- To create a backdoor for remote access to the infected system.
- To spread spam emails from the infected system.

Explanation:

Crypto-ransomware encrypts files on the infected system and then demands a ransom, typically in a form of cryptocurrency like Bitcoin, in exchange for the decryption key.

Crypto-ransomware does not primarily aim to steal sensitive information. It focuses on encrypting files and demanding a ransom. However, some advanced forms of ransomware may also steal data.

While some forms of malware are designed to create a backdoor for remote access, this is not the primary function of crypto-ransomware. Its main goal is to encrypt files and demand a ransom.

Spreading spam emails is a common function of some types of malware, but it is not the primary function of crypto-ransomware.

q_mal_fileless_secp8

Which virus operates only in memory and usually exploits a trusted application like PowerShell to circumvent traditional endpoint security solutions?

Answers:

- Remote Access Trojan (RAT)
- ***Fileless malware**
- Ransomware
- Worm

Explanation:

Fileless malware operates only in memory to avoid detection by traditional endpoint security solutions that are focused on matching signatures to files that have been written to the hard drive.

A worm is a self-replicating program.

Ransomware denies access to a computer system until the user pays a ransom.

A Remote Access Trojan (RAT) is a malware program that includes a backdoor that allows administrative control over the target computer.

q_mal_logic_bomb_secp8

What differentiates a logic bomb from other types of malware?

Answers:

- Logic bombs are always visible to the user.
- ***Logic bombs can only be triggered by a specific event or condition.**
- Logic bombs are always associated with ransomware.
- Logic bombs can only infect networked systems.

Explanation:

Logic bombs are unique in that they are programmed to execute (or "explode") when a specific condition or event occurs, such as a certain date and time, a specific user logging in, or a particular application being launched.

Logic bombs, like many other types of malware, often operate in the background without the user's knowledge until they are triggered.

While it's possible for a logic bomb to be part of a ransomware attack, they are not inherently associated with ransomware. Logic bombs can be used in various types of malicious software.

Logic bombs can infect any system, not just networked ones. They can be placed in any piece of software and triggered by a wide range of events or conditions.

q_mal_logic_secp8

Which of the following describes a logic bomb?

Answers:

- ***A program that performs a malicious activity at a specific time or after a triggering event.**

- A type of malicious code similar to a virus whose primary purpose is to duplicate itself and spread while not necessarily intentionally damaging or destroying resources.
- A program that appears to be a legitimate application, utility, game, or screensaver that performs malicious activities surreptitiously.
- A program that has no useful purpose but attempts to spread itself to other systems and often damages resources on the systems where it is found.

Explanation:

A logic bomb is a program that performs a malicious activity at a specific time or after a triggering event. Logic bombs can be planted by a virus, a Trojan horse, or an intruder. Logic bombs may perform their malicious activity at a specific time and date or when a specific event occurs on the system, such as logging in, accessing an online bank account, or encrypting a file.

A type of malicious code similar to a virus whose primary purpose is to duplicate itself and spread, while not necessarily intentionally damaging or destroying resources, is a worm.

A program that appears to be a legitimate application, utility, game, or screensaver that performs malicious activities surreptitiously is a Trojan horse.

A program that has no useful purpose but attempts to spread itself to other systems and often damages resources on the systems where it is found is a virus.

q_mal_ransomwae_secp8

Which of the following is a type of malware that prevents the system from being used until the victim pays the attacker money?

Answers:

- Remote Access Trojan (RAT)
- Denial-of-service attack (DoS attack)
- ***Ransomware**
- Fileless virus

Explanation:

A type of malware used to prevent the system from being used until a ransom is paid by the victim is known as ransomware.

While it does perform a denial-of-service, a DoS attack does not necessarily demand payment.

A Remote Access Trojan (RAT) is a malware program that includes a backdoor that allows administrative control over the target computer.

A fileless virus uses legitimate programs to infect a computer.

q_mal_rat_secp8

Which kind of malware provides an attacker with administrative control over a target computer through a backdoor?

Answers:

- Potentially unwanted program (PUP)

- Crypto-malware
- Trojan horse
- ***Remote Access Trojan (RAT)**

Explanation:

A Remote Access Trojan (RAT) provides a backdoor for an attacker to remotely control a computer with administrative control. The other types of malware could be used in conjunction with a RAT, but they do not provide remote control access.

PUP is software that contains adware, installs toolbars, or has other unclear objectives.

Crypto-malware is ransomware that encrypts files until a ransom is paid.

A Trojan horse is a malicious program that is disguised as legitimate or desirable software.

q_mal_rootkit_secp8

Which of the following are characteristics of a rootkit? (Select two.)

Answers:

- ***Requires administrator-level privileges for installation.**
- ***Resides below regular antivirus software detection.**
- Monitors user actions and opens pop-ups based on user preferences.
- Uses cookies saved on the hard drive to track user preferences.
- Collects various types of personal information.

Explanation:

A rootkit is a set of programs that allow attackers to maintain hidden, permanent, administrator-level access to a computer. A rootkit:

- Is almost invisible software.
- Resides below regular antivirus software detection.
- Requires administrator privileges for installation and then maintains those privileges to allow subsequent access.
- Might not be malicious.
- Often replaces operating system files with alternate versions that allow hidden access.

Spyware collects various types of personal information, such as internet surfing habits and passwords, and sends the information back to its originating source.

Adware monitors actions that denote personal preferences and then sends pop-ups and ads that match those preferences.

Both spyware and adware can use cookies to collect and report a user's activities.

q_mal_spyware_secp8

Which of the following BEST describes spyware?

Answers:

- ***It monitors the actions you take on your machine and sends the information back to its originating source.**
- It monitors user actions that denote personal preferences and then sends pop-ups and ads to the user that match their tastes.
- It is a program that attempts to damage a computer system and replicate itself to other computer systems.
- It is a malicious program disguised as legitimate software.

Explanation:

Spyware monitors the actions you take on your machine and sends the information back to its originating source.

Adware monitors the actions of the user that denote their personal preferences and then sends pop-ups and ads to the user that match their tastes.

A virus is a program that attempts to damage a computer system and replicate itself in other computer systems.

A Trojan horse is a malicious program disguised as legitimate software.

q_mal_worms_secp8

In 2001, a worm exploited vulnerabilities in Microsoft Internet Information Services (IIS) to infect over 250,000 systems in under nine hours.

What was this worm called?

Answers:

- Melissa
- ***Code Red**
- Nimda
- Michelangelo

Explanation:

The worm known as Code Red replicated across the internet with incredible speed using a vulnerability in Microsoft IIS.

In 1991, the Michelangelo virus was designed to infect MS-DOS systems and remain dormant until March 6, the birthday of Renaissance artist Michelangelo.

In 1999, the Melissa worm was the first widely distributed macro virus that was propagated in the form of an email message containing an infected Word document as an attachment.

In 2001, the Nimda worm took advantage of weaknesses found in the Windows platform and propagated itself in several ways, including email, infected websites, and network shares.

q_mal_prot_Anti-virus_secp8

You have installed antivirus software on the computers on your network. You update the definition and engine files and configure the software to update those files every day.

What else should you do to protect your systems from malware? (Select two.)

Answers:

- ***Educate users about malware.**
- ***Schedule regular full-system scans.**
- Enable account lockout.
- Disable UAC.
- Enable chassis intrusion detection.

Explanation:

You should schedule regular full-system scans to look for any malware. In addition, educate users about the dangers of downloading software and the importance of anti-malware protections.

You should enable User Account Control (UAC) to prevent unauthorized administrative changes to your system.

Use account lockout to help protect your system from hackers trying to guess passwords.

Use chassis intrusion detection to identify when the system case has been opened.

q_mal_prot_reimage_secp8

Which of the following malware recovery techniques is often faster and more effective than malware removal and cleanup of an infected computer?

Answers:

- ***Re-imaging the computer**
- Enabling privacy controls
- Remove removable drives
- Block specific executable files

Explanation:

Re-image a machine if your organization uses imaging solutions. Re-imaging or installing from scratch is often faster and more effective than malware removal and cleanup.

The following are preventative measures, not recovery techniques:

- Enabling privacy controls in Windows Internet Explorer such as deleting browser history.
- Removing removable drives to prevent unauthorized software from being installed on a system.
- Block executable files that have been copied from another computer.

q_mal_prot_web_filters_secp8

To prevent malware infection in your network system, you decide that it's critical to prevent malware attacks, such as ransomware and phishing, by restricting access to sites that might be malicious.

Which of the following BEST represents this type of prevention technique?

Answers:

- ***Web filters**

- Pop-up blocker
- Patching the operating system
- Updating your web browser

Explanation:

By installing web filters, you can prevent malware attacks, such as ransomware and phishing attacks, which often originate from malicious websites. By restricting access to these sites, web filters significantly reduce the risk of malware infections.

The following are also malware prevention techniques, but do not meet your prevention requirements:

- Using a pop-up blocker prevents adware.
- Installing the latest patches for the operating system.
- Updating your web browser

q_windefender_01

What is the current security intelligence version?

Answers:

- 1.229.426.0
- 1.125.333.1
- 1.229.426.0
- 1.300.000.1

q_windefender_02

After the update, what is the current security intelligence version?

Answers:

- 1.229.508.0
- 1.125.333.1
- 1.229.426.0
- 1.229.508.0
- 1.300.000.1

3.0 Cryptographic Solutions

3.1 Cryptography

As you study this section, answer the following questions:

- What is the difference between symmetric and asymmetric encryption?
- Which part of a simple cryptographic system must be kept secret—the cipher, the ciphertext, or the key?
- Which algorithms can be used to generate a hash?
- What is the process of digitally signing a message?
- What is a legitimate use for steganography?
- What are uses of blockchain in addition to cryptocurrency?
- What are the properties of a public/private key pair?

In this section, you will learn to:

- Use steganography to hide a file.
- Hide files with OpenStego.

Key terms for this section include the following:

Term	Definition
Cryptography	The science and practice of altering data to make it unintelligible to unauthorized parties.
Plaintext	Unencrypted data that is meant to be encrypted before it is transmitted, or the result of the decryption of encrypted data.
Ciphertext	Data that has been enciphered and cannot be read without the cipher key.
Algorithm	Operations that transform a plaintext into a ciphertext with cryptographic properties; also called a cipher.
Cryptanalysis	The science, art, and practice of breaking codes and ciphers.
Encryption	Scrambling the characters used in a message so that the message can be seen but not understood or modified unless it can be deciphered. Encryption provides for a secure means of transmitting data and authenticating users. It is also used to store data securely. Encryption uses different types of cipher and one or more keys. The size of the key is one factor in determining the strength of the encryption product.
Key	In cryptography, a specific piece of information that is used in conjunction with an algorithm to perform encryption and decryption.

Symmetric encryption	Two-way encryption scheme in which encryption and decryption are both performed by the same key. Also known as shared-key encryption.
Key length	Size of a cryptographic key in bits. Longer keys generally offer better security, but key lengths for different ciphers are not directly comparable.
Asymmetric algorithm	Cipher that uses public and private keys. The keys are mathematically linked, using either Rivel, Shamir, Adleman (RSA) or elliptic curve cryptography (ECC) alogrithms, but the private key is not derivable from the public one. An asymmetric key cannot reverse the operation it performs, so the public key cannot decrypt what it has encrypted, for example.
Public key	During asymmetric encryption, this key is freely distributed and can be used to perform the reverse encryption or decryption operation of the linked private key in the pair.
Private key	In asymmetric encryption, the private key is known only to the holder and is linked to, but not derivable from, a public key distributed to those with whom the holder wants to communicate securely. A private key can be used to encrypt data that can be decrypted by the linked public key or vice versa.
Blockchain	A concept in which an expanding list of transactional records listed in a public ledger is secured using cryptography.
Open public ledger	Distributed public record of transactions that underpins the integrity of blockchains.
Steganography	The practice of concealing a file, message, image, or video within another file, message, image, or video.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.2 Implement Encryption Technologies 4.2.1 Encrypt data communications

TestOut Security Pro	1.4 Explain the importance of using appropriate cryptographic solutions
	• Public key infrastructure (PKI)
	  ○ Public key
	  ○ Private key
	  ○ Key escrow
	• Encryption
	  ○ Asymmetric

	• Symmetric • Algorithms • Key length • Obfuscation • Steganography • Hashing • Salting • Digital signatures • Blockchain 2.4 Given a scenario, analyze indicators of malicious activity • Cryptographic attacks • Downgrade • Collision • Birthday

3.1.1 Cryptography Concepts (Lesson Video)

Transcript:

A security specialist's main goal is to keep data safe and out of the hands of hackers and other threats. To do that, you can lock down a system using multiple methods, but there's always a chance that a hacker will still get through. You can use cryptography to add an additional layer of protection and make it even more difficult for hackers to steal your data, even if they get inside.

In this video, we'll discuss what cryptography is, the basic concepts associated with cryptology, and some common cryptology methods.
Cryptography is defined as the process of writing or solving messages using a secret code. Cryptography isn't new. The earliest known use of cryptography was by the Egyptians in 1900 BC to deliver messages using encryption.
Encryption is the process of converting normal text into text that makes no sense at all, which is called ciphertext. A cipher is the method, or algorithm, used to encrypt the data. I know that might sound a little confusing, since these terms are often used interchangeably, so let's look at an example using a puzzle.
When you're completing a puzzle, you normally look at the completed picture on the box and then put the pieces together to match the picture. If you wanted to disguise the puzzle from prying eyes, you could use cryptography. In this case, you could encrypt the puzzle by flipping all the pieces over and scrambling them up, so no one can see the images printed on the puzzle pieces. Then, to allow you or anyone else to put the puzzle together, you'll need a method that's only given out to certain people. That's the cipher. We'll used a numbered cipher here. Using that cipher, anyone would be able to put the puzzle together.
One of the most common encryption ciphers is the Caesar cipher.
This cipher was created by Julius Caesar. When he wanted to sent a secret message, Caesar would shift each letter down the alphabet by three places. Then he could safely send the message to his commanders because anyone who intercepted the message couldn't read it anyway. Since only the commanders knew the cipher was to shift the letters up the alphabet by three, only they could read it.
Today, encryption is much more advanced than the Caesar cipher. You'll hear a lot about encryption keys, hashes, and digital signatures. To understand modern encryption, you need to understand these concepts.
Let's talk about encryption keys first. We install locks on our doors to keep our homes safe. These locks require a specific key. Encryption works using the same concept. When we want to encrypt data, we use an encryption key. Encryption keys are random strings of bits that are used to lock and unlock data. These keys are generated using ciphers, such as the Advanced Encryption Standard, or AES. They ensure each key is unique and difficult to crack.

There are two types of encryption methods, symmetric and asymmetric. Symmetric encryption uses the same key to encrypt and decrypt data. Asymmetric encryption uses one key to encrypt and then a different key to decrypt. These are referred to as public and private keys.
When you're working with these encryption methods, different techniques are used to help safeguard your data. One technique is known as hashing.
Hashing is the process of converting one value into another using a mathematical algorithm like MD5 or SHA. Hashes have several characteristics that make them very useful. For example, hashes are deterministic, meaning that the same data being hashed always results in the same hash value. Hashes can also be created very quickly and cannot be decrypted. Likewise, it's infeasible to find two messages with the same hash value. Even a small change to a message will change the hash value so extensively that the new hash value appears uncorrelated with the old hash value.
Because of these features, hashes are used in many applications.
A hash function is often used to verify passwords. Since storing or transmitting passwords in cleartext could result in a massive security breach, most systems only store the hash of the password. When a user needs to be authenticated, the password entered by the user is hashed and compared with the stored hash.
One of the problems with using hashed passwords is that several online sites have collected massive databases containing a hash for tens of million of possible passwords. Once a hash has been captured, it can be compared with the hashes found in the database, quickly resulting in the password used to create the hash. To keep this from happening, most hashes also incorporate something called a salt.
Salting the hash means adding a random number of characters to the password before the hash is created. For example, if this is the password to be hashed, a salt such as this may be added. This becomes the string to be hashed. Since the salt is randomly generated each time, even if the same password is used and can be varying lengths, it's virtually impossible to create a database containing all the possible salted passwords.
Hashes are also used to verify that a document hasn't been corrupted or tampered with by creating a hash for the document. When the same hash algorithm is calculated at the receiving end, the hashes should match. If not, then the data was either corrupted or tampered with during transmission and should not be trusted. In a similar fashion, hashes are also used to digitally verify a signature, like when you digitally sign an important document using the internet. In this case, only the signature is hashed instead of the entire document.
While there are many different cryptography methods that can be used to secure data, one of the most powerful methods used today is Elliptic Curve Cryptology, or ECC. Many websites use this method to secure connections and the data they send back and forth.
To truly understand ECC and how it works would require us to dive into some really complicated math that's beyond the scope of this lesson. For now, you just need to know that this method is the next generation of cryptology, and the algorithm used can generate smaller keys that are more secure than other cryptology methods.
The last encryption technique we'll discuss is call steganography. Steganography is the technique of hiding or concealing a file, message, image, or video within another file, message, image, or video. If you've ever used invisible ink to write a hidden message, you've used steganography.
One method of steganography is to use a special program to hide a message in another file, such as this image. When the recipient receives the image, they can use the same software to find the hidden message. If the file is intercepted by a hacker, all they'll see is the image.
That's it for this lesson. In this lesson, we discussed what cryptography is and how encryption and ciphers are used. We covered the key concepts of cryptography, including encryption keys, hashes, and digital signatures.
And then we looked at a few cryptology methods, including hashing, elliptical curve cryptology, and steganography.

3.1.2 Cryptography Facts

The goal of all IT security specialists is to keep data safe. Hackers find ways to circumvent firewalls, IPS devices, and other security protocols put in place. Cryptography is one additional layer of defense that can be used to protect data.

This lesson covers the following topics:

- Cryptography overview
- Cryptography concepts
- Cryptography methods

Cryptography Overview

Cryptography is the process of writing or solving messages using a secret code. This is the opposite of security through obscurity. Security through obscurity means keeping something a secret by hiding it. This is considered impossible (or at least high risk) on a computer system. With cryptography, it does not matter if third parties know of the existence and location of the secret because they can never understand what it is without the means to decode it. A form of cryptography called encryption has been used throughout the ages, mainly to keep messages out of the hands of enemies. Encryption is the process of converting normal readable text into something unintelligible called ciphertext. A cipher is the method or algorithm used to encrypt or convert the data.

The following terminology is used to discuss cryptography:

- **Plaintext (or cleartext)** - is an unencrypted message.
- **Ciphertext** - is an encrypted message.
- **Algorithm** - is the process used to encrypt and decrypt a message.
- **Cryptanalysis** - is the art of cracking cryptographic systems.

One of the more popular forms of encryption is the Caesar cipher. This encryption method works by shifting each letter in the alphabet a certain number of spaces to the right or left. In the example below, the cipher is shifting to the right by three letters. A becomes D, B becomes E, C becomes F, etc.

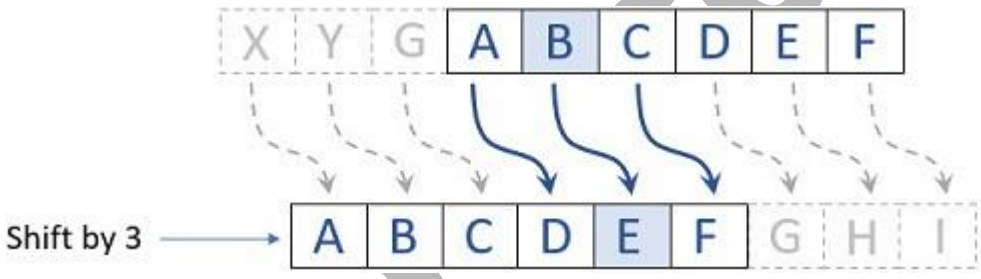

To decrypt the message, the reader must know how many spaces to shift the letters. For example, to encrypt the word TESTOUT using the Caesar Cipher with a shift of three to the right, TESTOUT would become WHVWRXW.

Below is the complete Caesar cipher using a shift of three letters to the right (you can see below which letters are used when they are shifted three spaces):

This is a very simple example of cryptography and is easily decrypted. With today's computing power, encryption methods used are much more complicated and powerful.

Cryptography Concepts

There are three main concepts to understand when dealing with today's encryption methods: encryption keys, hashing, and digital signatures.

Cryptography Concept	Description
Encryption keys	Encryption keys are used to encrypt and decrypt data. The key is a string of bits randomly generated using a specific cipher, such as Advanced Encryption Standard (AES). There are two types of encryption methods used with keys: symmetric and asymmetric. • Symmetric encryption uses the same key to encrypt and decrypt data. • Asymmetric encryption uses one key to encrypt the data and a different key to decrypt the data. These keys are known as a public key and private key.
Hashing	Hashing is the process of converting one value into another using a mathematical algorithm like MD5 or SHA. This fixed length of data is called the hash. • Hashing is used on data that does not need to be decrypted, such as passwords. • When a piece of data is run through a hashing algorithm, it always generates the same hash. If even one letter in a file has been altered, the resulting hash would be different. Because of this, hashing can be used to verify that data has not been altered during transmission. • A hash cannot be decrypted. However, when using hashing for passwords, many online sites have collected massive databases containing a hash for tens of millions (or more) of possible passwords. Once a hash has been captured, it can be compared with the hashes found in the database, quickly resulting in the password used to create the hash.
Salt	Salt, or salting the hash, means that a random number of characters are added to the password before the hash is created. For example, if the password to be hashed was p@ssw0rd, a salt, such as E1343135E119C253, may be added. Therefore, the string to be hashed would be p@ssw0rdE1343135E119C253. Since the salt is randomly generated each time, even if the same password is used and is of varying lengths, it's virtually impossible to create a database containing all the possible salted passwords.
Digital signatures	By combining a user's private encryption key and a hash of the data, a user can create a digital signature. A digital signature verifies that the data is legitimate and provides non-repudiation. This means that the sender cannot deny having sent the file.

Cryptography Methods

There are many different cryptography methods used today. One important thing to keep in mind is that all cryptography uses advanced math concepts to generate encryption keys and hashes.

Cryptography Method	Description
Elliptic Curve Cryptography (ECC)	Elliptic Curve Cryptography is one of the newer methods being implemented. ECC can generate smaller keys that are more secure than most other methods. Many websites today use ECC to secure connections and data transmissions.

Cryptography Method	Description
Perfect Forward Secrecy	This cryptography method is used quite often in messaging apps. Instead of the same key being used for an entire conversation or session on a website, each transmission is encrypted with a different unique key.
Steganography	Steganography is the technique of hiding or concealing a file, message, image, or video within another file, message, image, or video. Special programs are often used to hide messages in media files. If a hacker intercepts the message, all they see is the media. They don't know that there is a hidden message.

3.1.3 Symmetric vs Asymmetric Encryption (Lesson Video)

Transcript:

Keeping data and messages safe from prying eyes is the goal of cybersecurity. We can hide data from hackers by encrypting it. Encryption is the process of encoding data into a format that's unreadable. All encryption relies on a unique key to encrypt and decrypt the data. This key is basically a password that's combined with the cipher to encrypt the data. In this lesson, we're going to look at some encryption methods including symmetric encryption, asymmetric encryption, hybrid cryptosystems, and ephemeral keys.

Symmetric encryption uses the same key to encrypt and decrypt data. This is the simplest and oldest form of encryption. Let's look at how this works.
Travis has a confidential file that he needs to send to Craig. To ensure that the file is sent safely, Travis encrypts the file using his secret symmetric encryption key and sends it to Craig. To open the file, Craig must use the same encryption key to decrypt the file.
One of the biggest problems with symmetric encryption is that both parties must use the same key. If Travis and Craig send encrypted data often, they've probably set up this secret key in advance. But if not, Travis needs to let Craig know what the secret key is. Sending the key with the data defeats the purpose because if a hacker intercepts the message, they can decrypt the file.
Imagine what a pain it would be to distribute the secret key if multiple people were to need access to it. Not only is this difficult to do, but the more people that have the key, the more likely it is for the key to be compromised.
Symmetric encryption is very secure and works extremely well when you need to encrypt a large amount of data because it requires less CPU power than asymmetric encryption.
There are many symmetric-key algorithms used, but some of the more common ones are Data Encryption Standard, Rivest Cipher, and Advanced Encryption Standard.
The Advanced Encryption Standard, or AES, has for the most part replaced all other types of symmetric encryption. AES is a subset of the Rijndael block cipher, It was developed in 2001 during a competition run by the US National Institute of Standards and Technology, or NIST. AES offers three key lengths of 128, 192, or 256 bits. AES is used in all sorts of ways including 802.11 communication, BitLocker, and even some games engines.
Some other symmetric algorithms you may come across are Blowfish, IDEA, CAST, and Twofish. While these algorithms aren't used much anymore, you should at least know they exist.
Asymmetric encryption makes use of two different keys, a public key and a private key. The public key is used to encrypt the data and the private key is used to decrypt it. Since these keys mirror each other, the private key only decrypts data that was encrypted with the matching public key. Let's see this in action.
Travis is sending his confidential file to Craig. To do that, Travis must first get Craig's public key. Travis then uses Craig's public key to encrypt the data, and the file is sent to Craig. When Craig gets the file, he uses his private key to decrypt it.
As another example, asymmetric encryption is used in most communication over the internet.
For example, when you log on to a website, your browser sends a request for the public key which is then sent by the web server. Your login information is entered and encrypted using that public key, which is then sent to the server. The server in turn decrypts the login information using its private key. You're now logged into the web server. As you can see

from these examples, users don't ever reveal their private keys in this exchange. Data security is increased with asymmetric encryption, which is the primary reason for using it.
Another reason for using asymmetric encryption is to create digital signatures.
A digital signature is a method used to ensure that a received message or document hasn't been modified from the time it was signed.
This is done by generating a unique hash of the message or document and encrypting it with the sender's private key. When the recipient receives the message, they can verify the digital signature to ensure that the data or message is legitimate. Since anyone can decrypt these messages or documents with the public key, digital certificates don't secure data. They simply verify that the data is legitimate and came from the proper source. Guaranteeing the data source in this way provides non-repudiation, or the ability to ensure that the document's sender cannot deny sending the document.
There are four main asymmetric encryption algorithms you need to be aware of. These are Diffie-Hellman, RSA, DSA, and ECC.
Diffie-Hellman was the first widely used asymmetric algorithm. It was released in 1976 by Whitfield Diffie and Martin Hellman. This algorithm allows two parties who've never met to safely create a shared key over a public channel, such as the internet. We can see how this key is created using a choice of colors.
For example, our two users each agree on a paint color such as yellow. Each user then decides on a secret color to mix with their common color, creating a color they can share publicly. They send each other their results and then mix in their secret color. Each user ends up with the same new secret color. This new secret color, or key, is then used to encrypt and decrypt their messages. If a hacker were to intercept any of these messages without knowing each user's secret color, they wouldn't be able to reverse the process and figure out the final secret color. Diffie-Hellman is frequently implemented in security protocols such as TLS, IPSec, SSH, and more.
In 1977, Ron Rivest, Adi Shamir, and Leonard Adleman released one of the first public-key cryptosystems, which is known by the initial letters of the developer's surnames RSA. RSA defined the process of using a public key to encrypt data and a secret key to decrypt it. RSA is still one of the most widely used algorithms for securing data transmissions and creating digital signatures.
The Digital Signature Algorithm, or DSA, is a Federal Information Processing Standard for digital signatures. It was proposed in 1991 by NIST and became the government standard in 1993. It uses a different algorithm than RSA but provides the same level of security.
Elliptic Curve Cryptography, or ECC, is the latest encryption method. The algorithms used with this method can generate smaller keys that are just as secure as other methods. Many websites have adopted ECC to secure data being sent back and forth.
Asymmetric encryption methods aren't very efficient because they rely on some complicated mathematical computations. But they do make it easier to exchange keys. In contrast, Symmetric methods are efficient, but they become inconvenient when the time comes to get the keys out. A hybrid cryptosystem combines the efficiency of symmetric methods and the convenience of asymmetric methods.
Let's see how it works.
Travis needs to send a confidential file to Craig. Before doing so, Travis uses his symmetric private key to encrypt the data. Next, Travis encrypts the secret key using Craig's public key. The file is then sent to Craig. When Craig gets the message, he first uses his private key to decrypt the symmetric secret key and then uses the symmetric key to decrypt the message.
As long as Craig's private key is kept secret, the data is secure. Most secure communication methods such as TLS utilize a hybrid cryptosystem nowadays.
Generally, when we establish a session and exchange keys, all communication during that session uses the same encryption keys. These are known as static keys. The problem is that the longer a key is used, the more susceptible it is to attack. To combat this, we use ephemeral keys.
These are keys that are generated for each new session or transaction. Perfect forward secrecy makes use of ephemeral keys. Perfect forward secrecy and ephemeral keys are used quite often in instant messaging apps.
Every message that's sent generates its own encryption key. It doesn't matter if a hacker intercepts any of the keys because the next message has a completely different one. The hacker is unable to eavesdrop on the entire conversation or session.
That's it for this lesson. In this lesson, we first looked at symmetric encryption which uses the same key to encrypt and decrypt data. We then looked at asymmetric encryption which uses a public key and a private key to encrypt and decrypt data. We also covered hybrid cryptosystems that combine the efficiency of symmetric methods and the convenience of asymmetric methods. Finally, we covered ephemeral keys. These are keys that are generated for each new session or transaction.

3.1.4 Symmetric and Asymmetric Encryption Facts

There are three main types of cryptographic algorithms with different roles to play in the assurance of the security properties: confidentiality, integrity, availability, and non-repudiation. These types are hashing algorithms and two types of encryption ciphers: symmetric and asymmetric.

This lesson covers the following topics:

- Symmetric encryption
- Key length
- Asymmetric encryption
- Hybrid cryptosystems
- Ephemeral keys

Symmetric Encryption

An encryption algorithm or cipher is a type of cryptographic process that encodes data so that it can be stored or transmitted securely and then decrypted only by its owner or its intended recipient. Using a key with the encryption cipher ensures that decryption can only be performed by an authorized person.

Substitution and Transposition Algorithms

It is helpful to consider simple substitution and transposition algorithms to understand how an algorithm works. A substitution cipher involves replacing characters or blocks in the plaintext with different ciphertexts. Simple substitution ciphers rotate or scramble letters of the alphabet. For example, ROT13 rotates each letter 13 places, so A becomes N, for instance. The ciphertext "Uryyb Jbeyq" can be decrypted as the plaintext "Hello World."

In contrast to substitution ciphers, the units in a transposition cipher stay the same in plaintext and ciphertext, but their order is changed according to some mechanism. Consider how the ciphertext "HLOOLELWRD" has been produced:

H L O O L

E L W R D

The letters are written as columns, and then the rows are concatenated to make the ciphertext.

Modern encryption algorithms use these basic techniques of substitution and transposition in complex ways that can defeat attempts at cryptanalysis.

Symmetric Algorithms

A symmetric algorithm is one in which encryption and decryption are both performed by the same secret key. The secret key must be kept known to authorized persons only. If the key is lost or stolen, the security is breached. Symmetric encryption is used for confidentiality. For example, Alice and Bob can share a confidential file in the following way:

- Alice and Bob meet to agree which cipher to use and a secret key value. They both record the value of the secret key, making sure that no one else can discover it.
- Alice encrypts a file using the cipher and key.
- Alice sends only the ciphertext to Bob over the network.
- Bob receives the ciphertext and is able to decrypt it by applying the same cipher with his copy of the secret key.

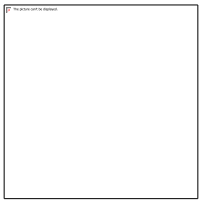

Symmetric encryption operation and weaknesses (Images © 123RF.com).

Symmetric encryption is very fast. It is used for bulk encryption of large amounts of data. The main problem is how Alice and Bob "meet" to agree upon or exchange the key. If Mallory intercepts the key and obtains the ciphertext, the security is broken.

Note that symmetric encryption cannot be used for authentication or integrity. Alice and Bob are able to create exactly the same secrets because they both know the same key.

The most common symmetric algorithm is the Advanced Encryption Standard (AES), also known as the Rijndael cipher. It was developed by Jaon Daemen and Vincent Rijmen in 2001 as part of a NIST competition held to find a replacement for DES. AES offers three different key lengths of 128, 192, and 256 bits. AES is used in many applications, including 802.11 communications, Bitlocker, and game engines.

Other common symmetric algorithms include Data Encryption Standard (DES), Rivest's Cipher (RC), International Data Encryption Algorithm (IDEA), Blowfish, Twofish, and CAST.

Key length

Encryption algorithms use a key to increase the security of the process. For example, if you consider the substitution cipher ROT13, you should realize that the key is 13. You could use 17 to achieve a different ciphertext from the same method. The key is important because it means that even if the cipher method is known, a message still cannot be decrypted without knowledge of the specific key.

A keyspace is the range of values that the key could be. In the ROT13 example, the keyspace is 25 (ROT1. . . ROT25). Using ROT0 or ROT26 would result in ciphertext identical to the plaintext. Using a value greater than 26 to shift through the alphabet multiple times is equivalent to a key from the 1-25 range. ROT0 and ROT26+ are weak keys and should not be used.

Modern ciphers use large keyspaces where there are trillions of possible key values. This makes the key difficult to discover via brute force cryptanalysis. *Brute force cryptanalysis* means attempting decryption of the ciphertext with every possible key value and reading the result to determine if it is still gibberish or plaintext.

Keys for modern symmetric ciphers use a pseudorandomly generated number of bits. The number of bits is the key length . For example, the most commonly used symmetric cipher is the Advanced Encryption Standard (AES). This can be used with two key lengths. AES-128 uses a 128-bit key length. A bit can have one of two values (0 or 1), so the number of possible key values is two multiplied by itself a number of times equivalent to the key length. This is written as 2^{128}, where 2 is the base, and ^{128} is the exponent. AES-256 has a keyspace of 2^{256}. This keyspace is not twice as large as AES-128; it is many trillions of times bigger and consequently significantly more resistant to brute force attacks.

The drawback of using larger keys is that the computer must use more memory and processor cycles to perform encryption and decryption.

Asymmetric Encryption

In a symmetric encryption cipher, the same secret key is used to perform both encryption and decryption operations. With an asymmetric algorithm , encryption and decryption are performed by two different but related public and private keys in a key pair.

When a public key is used to encrypt a message, only the paired private key can decrypt the ciphertext. The public key cannot be used to decrypt the ciphertext. The keys are generated in a way that makes it impossible to derive the private key from the public key. This means that the key pair owner can distribute the public key to anyone they want to receive secure messages from:

- Bob generates a key pair and keeps the private key secret.
- Bob publishes the public key. Alice wants to send Bob a confidential message, so they take a copy of Bob's public key.
- Alice uses Bob's public key to encrypt the message.
- Alice sends the ciphertext to Bob.
- Bob receives the message and is able to decrypt it using their private key.
- If Mallory has been snooping, they can intercept both the message and the public key.
- However, Mallory cannot use the public key to decrypt the message, so the system remains secure.

Asymmetric encryption (Images © 123RF.com)

The drawback of asymmetric encryption is that it involves substantial computing overhead compared to symmetric encryption. Where a large amount of data is being encrypted on disk or transported over a network, asymmetric encryption is inefficient. Rather than being used to encrypt the bulk data directly, the public key cipher can be used to encrypt a symmetric secret key. This allows Alice and Bob to exchange a bulk encryption session key without Mallory being able to learn it.

Asymmetric encryption can be implemented using a number of algorithms. Each algorithm has a different recommended key length. The Rivest, Shamir, Adelman (RSA) asymmetric cipher requires a 2,048-bit private key to achieve an acceptable level of security. The Elliptic Curve Cryptography (ECC) asymmetric cipher can use 256-bit private keys to achieve a level of security equivalent to a 3,072-bit RSA key.

The following table shows common asymmetric encryption algorithms:

Asymmetric Algorithm	Description
Rivest-Shamir-Adleman (RSA)	RSA was developed by Ron Rivest, Adi Shamir, and Leonard Adleman. RSA was released shortly after Diffie-Hellman in 1977. • RSA is still one of the most commonly used algorithms and helped define the process of using a public key to encrypt data and a private key to decrypt the data. • RSA is used extensively for creating digital signatures.
Elliptic Curve Cryptography (ECC)	Elliptic Curve Cryptology is one of the newer methods being implemented. It was originally introduced in 1985. It did not enter wide usage until 2004. • ECC is able to generate smaller keys that are more secure than most other methods. • Many websites today use ECC to secure connections and data transmissions.

Asymmetric Algorithm	Description
Diffie-Hellman	Released in 1976 by Whitfield Diffie and Martin Hellman. Its purpose was to allow two users who have never met to safely create a shared key over a public channel such as the internet. • Diffie-Hellman is used as follows: 1. The two users agree on two numbers: a prime number (P) and a generator (g). These numbers can be shared publicly. 2. Each user then randomly generates a private number, or key, unique to themselves. 3. Using the prime number, generator, and private key, each user generates a public key using the following formula: ▪ (G^ private number) MOD P 4. The users exchange their public keys, which are then used to create a shared secret key using the following formula: ▪ (Shared Public Key^ private number) MOD P 5. Because each public key was generated using the same prime number and generator, each user will come up with the same number for the shared secret key. 6. If a hacker intercepted any of the exchanges, they could not reverse the process without knowing each user's secret number. • Diffie-Hellman is frequently implemented in security protocols such as TLS, IPSec, SSH, and others.
Digital Signature Algorithm (DSA)	DSA was proposed in 1991 by NIST and became the government standard in 1993 • DSA is only used for creating digital signatures. • It uses a different algorithm than RSA but provides the same level of security.

Hybrid Cryptosystems

Hybrid cryptosystems combine the efficiency of symmetric encryption with the convenience of asymmetric encryption. A hybrid cryptosystem is used as follows:

1. User1 uses their symmetric private key to encrypt some data.
2. User1 then encrypts that symmetric private key using the recipient's public key and sends both to the recipient.
3. User2, the recipient, uses their private key to decrypt User 1's private key, which is then used to decrypt the message.
4. As long as User2's private key is kept secret, the data remains secure.

Hybrid cryptosystems are used with many secure communication methods today, such as TLS.

Ephemeral Keys

In traditional encrypted communications, static keys are used. This means that the same key is used throughout an entire session. The problem with this is that the longer the keys are used, the more susceptible they become to an attack. Ephemeral keys can be used to resolve this issue.

Ephemeral keys are generated for each new session or message sent. For example, perfect forward secrecy (PFC) uses ephemeral keys.

Some popular instant messaging apps make use of ephemeral keys to encrypt messages. Each message sent uses a unique key to encrypt it. If a hacker intercepts one key, the rest of the messages are still safe.

3.1.5 Cryptography Algorithm (Lesson Video)

Transcript:

In this lesson, we're going to take a deeper dive into the world of cryptography. The cornerstone of cryptography is the algorithm, or cipher. Cipher is just a fancy word for the series of steps that are taken for data to be encrypted or decrypted. In this video, we're going to look at several types of ciphers and how they work.

One of the first cipher types is the stream cipher. A stream cipher is a symmetric encryption method that encrypts 1 bit of plaintext at a time. This method is based on the one-time pad, or OTP, which was very popular during World War 2. OTP uses a symmetric encryption key that should be as long or even longer than the data being encrypted.
For example, if you had a small message that was only 15 bits long, a random 15-bit key would need to be generated. That key can then be used to encrypt your message using a process called Exclusive-ORing, or XORing. XORing means that we compare two strings, our message and our random key, to generate an output. We do this by lining them both up. If the bits match, a 0 is generated. If they don't match, a 1 is generated. The output is our ciphertext!
What makes this so interesting is that this demonstrates perfect secrecy. This means that, to a hacker, this ciphertext looks completely random. It would be just as difficult to brute-force the key as it would be to brute-force the data itself.
OTP only works properly if our secret key is only used once. If we use it more than once, a hacker could then start to decode the secret key. OTP works great if we're dealing with small messages. But what if we needed to encrypt every video in this course?
A video can range from 500 megabytes to up to 5 gigs or more. Since every video file would need to generate a key that's the same size or larger than the video, we would quickly end up running out of hard drive space. So with today's data types, the one-time pad just isn't very practical.
Stream ciphers offer a solution to this problem.
A stream cipher can use a smaller fixed-length key that can be used repeatedly throughout the encryption process. This key is called a seed key, and in our example we'll say that it's 2048 bits. This key is run through a pseudorandom number generator which outputs a new and unique encryption key called a keystream. Notice that this new key matches the size of our data. The keystream is then XORed with the data to give us our ciphertext.
While stream ciphers are less secure than the one-time pad, the advantage to this method is that it can be used in real time. One of the most widely used stream ciphers was RC4. But due to multiple vulnerabilities discovered, it's no longer in use.
Many symmetric encryption algorithms use the block cipher method. With block ciphers, instead of encrypting our data one bit at a time, the data is encrypted one chunk or block at a time. The common block sizes used are 64, 128, or 256 bits in length.
For example, let's say that we want to encrypt a file that's 800 bits in size. Using a 256-bit block size, the data will be encrypted 256 bits, or 32 bytes, at a time until the process is complete. Since the last block ends up being smaller than 256 bits, the algorithm pads the block with some random data to bring it up to 256 bits.
As shown here, there are several block cipher operation modes that can be utilized depending on the application or use. Let's talk about a few of them.
Electronic Code Book Mode, or ECB, is the simplest block cipher operation mode.
With this mode, each block of plaintext data is encrypted separately. Several blocks can be encrypted at the same time, allowing for faster encryption. The biggest disadvantage is that identical data results in the same ciphertext, meaning that a hacker would be able to decipher some of the data if they intercept multiple blocks.
Cipher Block Chaining is like ECB, except that it uses an initialization vector. The initialization vector is a starting variable that's XORed with the current block's plaintext to encrypt the data. The first initialization vector is a random string and

each subsequent initialization vector is the ciphertext from the previous block. CBC is more secure than ECB but slower because multiple blocks can't be encrypted at the same time.
Cipher Feedback Mode, or CFB, also utilizes an initialization vector. But instead of using it on the plaintext, the initialization vector is encrypted and then that result is XORed with the plaintext to create the ciphertext block. This is the equivalent of encrypting the plaintext with a one-time pad.
Output Feedback Mode, or OFB, is practically identical to Cipher Feedback Mode. The difference is the initialization vector used after the first round of encryption. The CFB's output is XORed with the plaintext and the result is the next block's initialization vector. The process is the same in Output Feedback Mode, except that the encryption's output is the next block's initialization vector before it's XORed with the plaintext.
Next we have Counter Mode. Similar to ECB, every encryption process in Counter Mode is separate. Instead of using an initialization vector, Counter Mode uses a nonce combined with an encrypted counter. Nonce is simply a fancy word for a random string that's used for all blocks. The nonce's encrypted output and counter are then XORed with the plaintext to create the ciphertext. The counter starts at 0 and increments every block. By combining the nonce and counter, each block is using a different value so that even if the data is the same, the output will be different.
If we want to provide authentication with confidentiality, we can use the Galois Counter Mode, or GCM. This mode works just like Counter Mode. First, we have the counter and nonce encrypted. The output is XOR-ed with the plaintext to give us the ciphertext. Here's where things get different. The ciphertext is combined with a special hash. That output contains the ciphertext along with a message authentication code that gives us assurance that the data hasn't been tampered with. Because this method is efficient and provides authentication, we see it used often with network communications such as wireless networks and web servers that use SSH or TLS.
The other block cipher modes are classified as unauthenticated encryption and GCM is classified as authenticated encryption. There are some other encryption methods that authenticate, but the most common one is the Galois Counter Mode.
We're putting more and more Internet of Things, or IoT, devices into our networks; the need to encrypt data has never been more important. In 2018, the National Institute of Standards and Technology, or NIST, began the process of standardizing encryption algorithms called lightweight cryptography for these types of devices.
Many IoT devices are small, low-powered devices that don't have the resources to handle the encryption methods we've covered. IoT devices have limitations such as only having a small amount of RAM and CPU power. Some devices even run on batteries. Lightweight cryptography algorithms must take this into account and still provide high levels of security while being efficient.
One huge risk with all encryption is that the data must first be decrypted before it can be used. Homomorphic encryption addresses this concern by allowing encrypted data to be used without decrypting it first. This is done by performing mathematical operations on the ciphertext instead of on the actual data itself. Homomorphic encryption works like other asymmetric encryption methods by using a public key to encrypt the data and a private key to decrypt it.
It helps to look at data as math problems to better understand the differences in the types of homomorphic encryption. There are three types of homomorphic encryption: Partially homomorphic, which allows for only simple mathematical operations like addition and subtraction to be performed; somewhat homomorphic, which allows for more complex math such as multiplication to occur but only for a limited number of times; and fully homomorphic, which can handle any mathematical operation an unlimited number of times. Homomorphic Encryption has the potential to be a game changer, but as of now it's incredibly slow and inefficient.
That's it for this lesson. In this lesson, we looked at the different ciphers used in cryptography. We first examined the stream cipher which encrypts data one bit at a time. Then we looked at how block ciphers work and the different operation modes. Finally, we looked at two newer forms of cryptography. Lightweight cryptography is being designed for IoT devices and homomorphic encryption allows data to be used without having to decrypt it.

3.1.6 Cryptography Algorithms Facts

The cornerstone of all cryptography is the algorithm or cipher. There are different types of ciphers in use today.

This lesson covers the following topics:

- Stream cipher
- Block cipher
- Lightweight cryptography
- Homomorphic encryption

Stream Cipher

A stream cipher is a symmetric encryption method that encrypts data one bit at a time. The stream cipher is based on the one-time pad (OTP) concept used extensively during World War 2. Consider the following OTP characteristics:

- The OTP uses a symmetric encryption key the same length as the data being encrypted.
- The encryption key is used to encrypt the data using a process called XORing. This means that two binary strings are compared. If the bits match, a zero is generated. If the bits do not match, a one is generated. The output of the XOR process is the ciphertext.
- The OTP demonstrates what is called perfect secrecy. This means that it is just as, if not more, difficult for a hacker to brute-force the key as it would be to brute-force the data itself.
- The OTP only provides perfect secrecy if the secret key is only used once. If used more than once, a hacker could begin to decode it.
- OTP works well when dealing with small messages, but it becomes impractical when dealing with large data due to the keys being so extensive.

A stream cipher solves the problem of having excessively large keys associated with OTP by using a smaller, fixed-length seed key, such as one that is 2048 bits in length. The following describes the process of creating a seed key:

1. The seed key is run through a pseudorandom number generator, which outputs a new and unique encryption key the same size as the data being encrypted. This new key is called a keystream.
2. The keystream is XORed with the data to create the ciphertext.
3. The seed key can be used repeatedly throughout the encryption process.

While stream ciphers are less secure than the one-time pad, this method can be used in real-time. One of the most widely used stream ciphers was Rivest's Cipher 4 (RC4). However, multiple vulnerabilities have been discovered in RC4, rendering it insecure. It is no longer used today.

Block Cipher

Many symmetric encryption algorithms use the block cipher method. Instead of encrypting our data one bit at a time, a block cipher encrypts the data one chunk at a time. Common block sizes are 64, 128, or 256 bits in length. For example, when encrypting a piece of data that is 100 bytes in size using a 256-bit block size, the data would be encrypted 32 bytes at a time (256 bits/8 = 32 bytes). Because the last block is less than 32 bytes, extra random bits are added to bring the block up to 32 bytes.

There are six block cipher modes of operation that can be utilized depending on the application or use:

Mode of Operation	Description
Electronic Code Book (ECB)	ECB is the simplest mode of operation. • Each block of plaintext data is encrypted separately. • Blocks of data can be encrypted simultaneously, allowing for faster encryption. • The biggest disadvantage is that blocks with identical data will generate the same ciphertext.
Cipher Block Chaining (CBC)	CBC is similar to ECB, except this mode uses an initialization vector (IV).

Mode of Operation	Description
	• The IV is a starting variable that is XORed with the plaintext of the current block to encrypt the data. • The IV for the starting block is a randomly generated value. Each subsequent IV is the ciphertext from the previous block. • CBC is more secure than ECB due to the IV, but it is slower because blocks cannot be encrypted simultaneously.
Cipher Feedback mode (CFB)	CFB also uses an IV, but instead of using it on the plaintext, the IV is encrypted first. That output is then XORed with the plaintext to create the block of ciphertext. • This is the equivalent of using a one-time pad to encrypt the data. • The IV for the starting block is a randomly generated value. Each subsequent IV is the ciphertext from the previous block.
Output Feedback mode (OFB)	This mode is identical to CFB except for the IV used after the first round. • The output of the IV encryption is used as the next block's ciphertext.
Counter mode (CTR)	Instead of using an initialization vector, CTR uses a nonce combined with a counter that is encrypted. • A nonce is a random string used for all blocks during the encryption process. • The encrypted output of the nonce and counter is then XORed with the plaintext to create the ciphertext. • The counter increments for each block. This ensures that each block uses a different value so that even if blocks have the same data, the ciphertext will be different.
Galois Counter mode (GCM)	All other modes of operation are unauthenticated forms of encryption. The Galois Counter mode provides both encryption and authentication. • GCM works like Counter mode, except the ciphertext is combined with a special hash. • The output of the ciphertext and hash contains the encrypted data and a Message Authentication Code (MAC) that gives assurance that the message has not been tampered with. • Because GCM is extremely efficient and provides authentication, it is often used with network communications such as 802.11 and when sending encrypted data to a web server using TLS or SSH. • There are other encryption methods that also provide authentication, but GCM is the most widely used method.

Lightweight Cryptography

In 2018, NIST began the process to standardize encryption algorithms called lightweight cryptography. Lightweight cryptography is meant to be used on Internet of Things (IoT) devices.

Many IoT devices are small, low-powered devices that do not have the resources to handle other encryption methods. Some of the limitations of IoT devices that lightweight cryptography needs to address are:

- Small amount of RAM
- Low CPU power
- Low-powered or runs on batteries

Lightweight cryptography algorithms need to work on these devices efficiently while still providing high levels of security.

Homomorphic Encryption

An inherent risk with all encryption is that for the data to be worked on (computation on ciphertexts), it must first be decrypted. Homomorphic encryption addresses this concern by allowing data to be worked on without decrypting it first.

There are three types of homomorphic encryption. To explain the differences, it works best to think of data as integers and to use math functions to represent the manipulation of that data.

Homomorphic Encryption Type	Description
Partially homomorphic encryption (PHE)	PHE allows only select simple math functions (such as addition) to be performed. This means that only one math function can be performed an unlimited number of times on the encrypted values.
Somewhat homomorphic encryption (SHE)	SHE allows more complex math (such as multiplication) to occur. However, it can only be performed a limited number of times.
Fully homomorphic encryption (FHE)	This method can handle both simple and advanced math functions (such as addition and multiplication) being performed an unlimited number of times on the encrypted values. FHE is still in the developmental stage.

3.1.7 Identify Cryptographic Modes of Operation

3.1.8 Blockchain (Lesson Video)

Transcript:

A blockchain is a unique and increasingly popular implementation of cryptography. Blockchain technology was developed in 2008 and really came to the public with the release of Bitcoin. It's a decentralized and distributed ledger that uses cryptography to keep data secure and records and verifies transactions between two parties. In this lesson, we'll look at how a blockchain works and some interesting implementations of this newer technology.

Bitcoin cryptocurrency is the first prevalent use of a blockchain.
The idea is to provide a way for people to make online transactions without the use of a centralized third party, such as a bank. Let's take a closer look at the blockchain and how it works.

Each block is a transaction that's stored in a public database, or the chain. The block contains information about the transaction, including time, date, parties involved, and a unique hash that separates it from other blocks on the chain.
Let's look at this blockchain in action.
Travis is purchasing something from Craig and needs to send him money. Using his personal secret key and Craig's public key, Travis requests the transaction, which is represented online as a block. The transaction is sent out on a peer-to-peer network that consists of a bunch of computers called nodes.
The network verifies the transaction using known algorithms. The transaction is added to the chain, which creates a permanent record that can't be altered or denied. Travis's ownership of the money now transfers to Craig. The entire transaction happens very quickly, usually within a matter of minutes.
What makes the blockchain so secure is that every node contains a copy of the ledger. If a hacker changes a transaction, it won't matter because a copy of the transaction is also on many other computers.
Not just anyone can join a blockchain network, however. Each node must prove themselves and actually do some work. This is done through a process called mining.
Mining is adding blocks to the chain. To do this, extremely complex math computations take place to generate a 64-digit hexadecimal hash. Creating it requires a lot of power and time, so the process is split between each node. As a reward for performing these tasks, the nodes are usually paid with cryptocurrency, like Bitcoin.
Even though each transaction is publicly visible, all personal information is encrypted and hidden from everyone, even the nodes. Because it's a shared ledger and every transaction is publicly available, all persons involved can be held accountable for their actions.
Sending money isn't the only way blockchain can be used. Some large companies, including Microsoft and IBM, are working to adapt different technologies to the blockchain.
Many transactions that currently rely on a paper-based system could benefit from blockchain. For example, car or house titles could be transferred over the blockchain. Because the transactions are transparent, this would create a clear picture of legal ownership.
We could also use the blockchain to track the movement of things like food from its origin to the market. Any product supply chain could be followed, and the fact that it's all transparent would allow everyone peace of mind because we could see every stop the product has made.
There are many other ways the blockchain could be used, and as the technology becomes more popular, we'll see the blockchain serve many functions.
That's it for this lesson. In this video, we discussed what a blockchain is and how it works. Then we looked at some implementations of blockchain technology.

3.1.9 Blockchain Facts

Blockchain is a unique and increasingly popular implementation of cryptography that was developed in 2008.

This lesson covers the following topics:

- Blockchain
- Blockchain transaction
- Blockchain implementations

Blockchain

Blockchain is a concept in which an expanding list of transactional records is secured using cryptography. Each record is referred to as a *block* and is run through a hash function. The hash value of the previous block in the chain is added to the hash calculation of the next block in the chain. This ensures that each successive block is cryptographically linked. Each block validates the hash of the previous block all the way through to the beginning of the chain, ensuring that each historical transaction has not been tampered with. In addition, each block typically includes a time stamp of one or more transactions as well as the data involved in the transactions themselves.

The blockchain is recorded in an **open public ledger** . This ledger does not exist as an individual file on a single computer; rather, one of the most important characteristics of a blockchain is that it is decentralized. The ledger is distributed across a peer-to-peer (P2P) network in order to mitigate the risks associated with having a single point of failure or compromise.

Blockchain users can, therefore, trust each other equally. Likewise, another defining quality of a blockchain is its openness—everyone has the same ability to view every transaction on a blockchain.

Blockchain technology has a variety of potential applications. It can ensure the integrity and transparency of financial transactions, legal contracts, copyright and intellectual property (IP) protection, online voting systems, identity management systems, and data storage.

The first big use of the blockchain was the cryptocurrency, Bitcoin. The purpose of Bitcoin is to provide a method for people to make online transactions without the use of a centralized third party such as a bank.

Blockchain Transaction

Each block contains information about the transaction, including:

- Time
- Date
- Parties involved
- A unique hash that separates the block from other blocks on the chain

Each block goes through the same process:

1. User1 requests a transaction with User2. The request is made using User1's personal secret key and User2's public key.
2. The transaction is represented online as a block.
3. The block is distributed to everyone on a peer-to-peer network.
4. The network users verify that the transaction is valid.
5. The block is added to the chain. This provides an indisputable and transparent record of the transaction.
6. The contents of the transaction move to User2.

The entire transaction happens very quickly, usually within a matter of minutes. Every node contains a copy of the ledger. If a hacker changed a transaction, it would not matter because a copy of the transaction is stored on many other computers.

Each node plays an important part in the blockchain. Each node must prove itself by performing work before being allowed to join the network. This is typically done through mining.

Mining is the process of adding blocks to the chain. Each block has a 64-digit hexadecimal hash generated by extremely complex math computations. Generating the hash requires a lot of time and computing power. The process is split between each node. As a reward for performing these computations, the nodes are usually paid with cryptocurrency, like Bitcoin.

Even though each transaction is publicly visible, all personal information is encrypted and hidden from everyone, even the nodes. Because it is a shared ledger and every transaction is publicly available, all persons involved can be held accountable for their actions.

Blockchain Implementations

Sending money is not the only way blockchain can be used. Some large companies, including Microsoft and IBM, are working to adapt different technologies to the blockchain.

Many transactions that currently rely on a paper-based system could benefit from blockchain. For example, a car or house title could be transferred over the blockchain. Because the transactions are transparent, this would create a clear picture of legal ownership.

A blockchain could be used to track the movement of a product, such as food. The ability to follow and see every stop the product has made on its way to the consumer would provide assurances of the product's safety.

3.1.10 Use Steganography to Hide a File (Demo Video)

Transcript:

Sometimes, in order to maintain access to a network or system, malicious users try to transfer files. In this demo, we'll discuss how to hide the contents of a file inside another file using a technique called steganography. People have used steganography to send secret messages for hundreds of years. Typically, this method is used when the sender is concerned that their traffic is being inspected by a third party.

We're going to use a free tool called OpenStego that's able to hide a file inside of an image. This tool is very easy to use and has a graphical user interface. We're going to run the tool on Linux, but it's also available on Windows.
When we open up the tool, the first thing we see is the Hide Data screen; since this is the screen we want, let's move on to the fields that need to be filled out. Message File is the name of the file we want to hide in the Cover File, and it will be saved to the Output Stego File. I already have a file I want to hide, so I'll select that now. I'll come over here and navigate down until I find my file, secret-information.txt. Let's click OK.
Now that we've selected the message file, we need to select the file we want to save the message into. I've already downloaded a picture, so we'll select it and then continue. Once again, I'll navigate to my Pictures folder, where I've stored the PNG we're going to use.
Lastly, we need to select where we would like to save the file that we've encoded information into. Let's just call it secret-forest.png and save it in the same folder. I'll click over here and enter in 'secret-forest.png'. We could enter a password here, but we'll just leave it for now. Since that's all the information we need, let's click Hide Data.
That's it! To see that it worked, let's check to see that the Secret Forest file has been created. I'll browse through my folders to get to the Pictures folder. This is the file we created that contains the hidden data. Let's take a look at the original file to see whether there are any major visual differences. This side-by-side comparison shows that, visually, the images appear to be identical. If we look at the file sizes, though, we can see a bit of a difference between the two. In the lower right-hand corner, you can see that the original image is 10 megabytes, and the file that we encoded data into is almost 12 megabytes. This kind of difference would be difficult to detect if you didn't know the size of the original file.
Now that we see that the new file has been created, we should make sure that the data was stored properly. Let's try to extract the data. Let's go back to OpenStego and click on the Extract Data button. This page is a bit simpler, since the program only needs the new file and a folder to save the extracted file. I'll provide that data now and attempt to extract the data. I'll come over here and select my file. Click OK. Next, we'll just save it in the Pictures folder and select the Extract Data button. We'll wait for a second while it processes.
You can see that a text file was created in the folder. The file has the same name as the file we originally stored in the image, since the steganography tool saves that information in the image as well. We should make sure that the data inside the file is actually correct by opening it. I'll come up here and open the text file. And here's all our secret passwords.
In this demo, we discussed steganography and methods for saving data to an image and extracting data from an image.

3.1.11 Hide Files with OpenStego (Simulation)

Scenario

You are the IT security administrator for a small corporate network. Recently, some of your firm's proprietary data leaked online. You have been asked to use steganography to encrypt data into a file to be shared with a business partner. The data will allow you to track the source if the information is leaked again.

In this lab, your task is to use OpenStego to hide data in photos as follows:

- Encrypt and password-protect the user data in the file to be shared.
    - Message file: **John.txt**

    - Cover file: **gear.png**
    - Output Sego file: **send.png** (saved in the Documents folder)
    - Password: **NoMor3L3@ks!**
- Confirm the functionality of the steganography by:
    - Extracting the data to **C:\Users\Administrator\Documents\Export** .
    - Open the extracted file to confirm that the associated username has been embedded into the file.

Explanation

Complete this lab as follows:

1. Encrypt the user data into the file to be shared.
    a. In the search field on the taskbar, type **OpenStego** .
    b. Under *Best match* , select **OpenStego** .
2. Select the Message, Cover, and Output Stego files.
    a. For the *Message File* field, select the **ellipses** [...] button at the end of the field.
    b. Double-click **John.txt** to select the file.
    c. For the *Cover File* field, select the **ellipses** [...] button at the end of the field.
    d. Double-click **gear.png** to select the file.
    e. For the *Output Stego File* field, select the **ellipses** [...] button at the end of the field.
    f. In the *File name* field, enter **send.png** and then select **Open** .
3. Password protect the file.
    a. In the *Password* field, enter **NoMor3L3@ks!**
    b. In the *Confirm Password* field, enter **NoMor3L3@ks!**
    c. Select **Hide Data** .
    d. Select **OK** .
4. Extract the data.
    a. Under *Data Hiding* , select **Extract Data** .
    b. For the *Input Stego File* field, select the **ellipses** [...] button.
    c. Double-click **send.png** to select the file with the encryption.
    d. For the *Output Folder for Message File* field, select the **ellipses** [...] button.
    e. Double-click **Export** to set it as the destination of the file output.
    f. Click **Select Folder** .
    g. In the *Password* field, enter **NoMor3L3@ks!** as the password.
    h. Select **Extract Data** .
    i. Select **OK** .
5. Verify that the decryption process was successful.
    a. From the taskbar, select **File Explorer** .
    b. Double-click **Documents** to navigate to the folder.
    c. Double-click **Export** to navigate to the folder.
    d. Double-click **John.txt** .

3.1.12 Cryptographic Attacks (Lesson Video)

Transcript:

No matter how many barriers we put in place to protect our data, hackers will always try to figure out a way to get around them and steal valuable information.

In this video, we're going to look at some common cryptographic attacks and what the future of cryptography might look like.

Encrypted data is usually secure, but there are no guarantees in the security world. By using new, targeted, or sophisticated attacks, hackers might be able to decrypt data you thought was safe. Some of the more common attacks are dictionary attacks, collision attacks, birthday attacks, and downgrade attacks.
A dictionary attack is a type of brute force attack that is used quite often. In this attack, the hacker has a huge list of words and phrases that are used to guess a decryption key. These attacks work well against weak passwords, such as password123. Using longer and uncommon passphrases is the key to beating dictionary attacks.
Another type of attack is the collision attack. A collision attack tries to find two inputs that produce the same hash value. This type of attack is often used on digital signatures.
For example, we have a hacker that creates two different documents that have an identical hash value. With this type of attach, the hacker wants to get Travis to accept the second document by making it seem that it came from Craig. The first document is sent to Craig, who agrees with what the document says, so he signs it and sends it back. The hacker then takes Craig's signature from that document, attaches it to a bad document, and sends that to Travis. Because the documents have the same hash, Travis' software is unable to detect that there's anything wrong, and it looks totally legitimate. The problem with this attack is the difficulty of generating two documents that have the same hash.
A birthday attack combines a collision attack and a brute force attack. This attack's name is taken from the birthday probability math problem.
If you have 30 people are in a room, the chance that someone has the same birthday as you is about 8 percent. But the probability that any two people in the room have the same birthday is 70 percent. This is because we're not looking for an exact match, just any match. This concept can be used to create two documents with the same hash, as in the collision attack example.
Another popular attack is the downgrade attack. Downgrade attacks force the system to use an older, less secure protocol for communication.
A common example of this attack is SSL exploitation. Many servers have both SSL and TLS installed on them. If a client can't use TLS, then the server can revert to SSL. A hacker can set up their computer to only use SSL, which will then allow them to launch SSL-based attacks on the server.
In this attack, the hacker intercepts the HTTPS packets and reverts them to HTTP packets. If the server isn't configured properly, it'll respond using HTTP, which allows the hacker to see everything being sent back and forth.
To combat these attacks, it's imperative that servers don't support older, less secure protocols. Proper server configuration is the key to stopping these types of attacks.
Now let's look ahead to what's coming in the security world. The future of computers and cryptography lies in quantum computing.
Classic computing works by processing bits, or 0s and 1s. These bits represent the electrical signals of on and off. Quantum computing uses qubits, which can exist as both a 1 and 0 at the same time. Quantum computing is so much more powerful than today's computing.
This increased computing power means that today's encryption standards can be hacked easily and quickly. An encryption key that might take years to crack with today's computers can take days or even hours with quantum computers.
To combat the inevitable increase of quantum computing, researchers have already started work on post-quantum cryptography. These new methods will be used to ensure the safety of our data.
That wraps up this lesson. In this video, we covered different types of cryptographic attacks, including dictionary, collision, birthday, and downgrade attacks. Then we looked at the future of cryptography, which involves quantum computing and post-quantum cryptography.

3.1.13 Cryptographic Attack Facts

Hackers attempt to figure out a way to get to data they want. Encrypting data is usually relatively secure, but there is, unfortunately, no such thing as a sure thing when it comes to protecting data. By using different types of attacks, hackers might be able to gain access to encrypted data.

This lesson covers the following topics:

- Common cryptographic attacks
- Future of cryptography

Common Cryptographic Attacks

The following table covers some of the more common cryptographic attacks.

Attack Method	Description
Dictionary	A dictionary attack is a type of brute-force attack. The hacker uses a list of words and phrases to try to guess the decryption key. • Dictionary attacks work well if weak passwords are used. • Using longer and uncommon passphrases is the best way to secure data against these attacks.
Collision attack	A collision attack tries to find two inputs that produce the same hash value. This type of attack is often used on digital signatures. • If a hacker wanted to get User2 to sign a document by making it seem like it came from User1, they would generate two documents that generate the same hash. • The hacker would send one document to User1 and get that signature. • The signature would be attached to the second document and sent to User2. Because the hashes are identical, User2 thinks the document is legitimate and has been signed by User1. Generating longer hash outputs is the key to stopping these types of attacks.
Birthday attack	This attack combines a collision attack and a brute-force attack. The name is taken from the birthday probability math problem. The birthday probability math problem states that if you have 30 people in a room, the probability that someone has the same birthday as you is approximately 8%. However, the probability that any two people in the room have the same birthday is 70%. This is because we are not looking for an exact match (just any match), so the probability is higher. Digital signatures can be susceptible to birthday attacks. Generating longer hash outputs is the key to stopping these attacks.
Downgrade attack	A downgrade attack forces the system to use an older, less secure communication protocol. • SSL exploitation is a common implementation of this attack. A hacker can set up their computer to only use SSL so that when the request is sent to the server, the server downgrades from TLS to SSL to communicate. This then allows the hacker to launch SSL-based attacks on the server. • Downgrade attacks are often used as part of a man-in-the-middle (MITM) attack. The hacker can intercept an HTTPS packet and downgrade it to an HTTP packet. If the server is not configured properly, the server responds using HTTP. This allows the hacker to now see all communications.

	To prevent downgrade attacks, servers must be set up not to support these older and less secure protocols. Proper server configuration is the key to stopping these types of attacks.

Future of Cryptography

The future of computers and cryptography lies in quantum computing. Classic computing works by processing bits of 1s and 0s. These bits represent electrical signals, on and off. Quantum computing uses qubits which can exist as both a 1 and 0 at the same time. Quantum computing is exponentially more powerful than today's computing standards.

This increased computing power means that today's encryption standards can be hacked easily and quickly. An encryption key that might take years to crack with today's computers can take days or even hours with quantum computers. To combat the inevitable increase of quantum computing, researchers have already started work on post-quantum cryptography. These new methods will be used to ensure the safety of our data in the future.

3.1.14 Practice Questions (Section Quiz)

q_cryp_concepts_algorithms_secp8

You are a cybersecurity analyst at a large corporation. The company has recently received a series of suspicious emails containing encrypted messages.

You suspect that the messages are using a combination of substitution and transposition algorithms for encryption.

The most recent message reads: "HLOOLELWRD".

Which of the following steps would you take to decrypt this message?

Answers:

- Apply a Caesar cipher with a shift of 3 to the right.
- Apply a ROT13 substitution cipher.
- ***Rearrange the letters in blocks of two.**
- Apply a Caesar cipher with a shift of 3 to the left.

Explanation:

Rearrange the letters in blocks of two is the correct answer. The message appears to be encrypted using a simple columnar transposition cipher, where letters are written in columns and then concatenated in rows. Rearranging the letters in blocks of two (i.e., treating "HL", "OO", "LE", "LW", "RD" as separate blocks) would result in the plaintext message "HELLO WORLD".

Apply a Caesar cipher with a shift of 3 to the right is incorrect. A Caesar cipher is a type of substitution cipher, not a transposition cipher. Applying a Caesar cipher to this message would not result in a meaningful decryption.

Apply a ROT13 substitution cipher is incorrect. ROT13 is a specific type of Caesar cipher that shifts letters 13 places. However, this does not account for the transposition aspect of the encryption.

Apply a Caesar cipher with a shift of 3 to the left is incorrect. A Caesar cipher is a type of substitution cipher and would not account for the transposition aspect of the encryption.

q_cryp_concepts_caesar_cipher_secp8

Which of the following encryption methods works by shifting each letter in the alphabet a certain number of spaces to the right or left?

Answers:

- ***Caesar cipher**
- Encryption key
- Hashing
- Salt

Explanation:

One of the more popular forms of encryption is the Caesar cipher. This encryption method works be shifting each letter in the alphabet a certain number of spaces to the right or left. An example of this cipher is shifting to the right by three letters: A becomes D, B becomes E, C becomes F, and so forth.

Encryption keys are used to encrypt and decrypt data. The key is a string of bits that is randomly generated using a specific cipher, such as Advanced Encryption Standard (AES). There are two types of encryption methods used with keys: symmetric and asymmetric.

Hashing is the process of converting one value into another using a mathematical algorithm like MD5 or SHA. This fixed length of data is called the hash.

Salt, or salting the hash, means that a random number of characters are added to a password before the hash is created.

q_cryp_concepts_ecc_secp8

You are creating a website for a financial investment company customers and are using a cryptography method that secures connections and data transmissions by generating smaller keys that are more secure than most other methods.

Which of the following cryptography methods are you using?

Answers:

- ***Elliptic Curve Cryptography (ECC)**
- Perfect Forward Secrecy
- Steganography
- Digital signature

Explanation:

You are using Elliptic Curve Cryptography (ECC). ECC is one of the newer methods being implemented. It is able to generate smaller keys that are more secure than most other methods. Many websites today use ECC to secure connections and data transmissions.

The Perfect Forward Secrecy cryptography method is used quite often in messaging apps. Instead of the same key being used for an entire conversation or session on a website, each transmission is encrypted with a different unique key.

Steganography is the technique of hiding or concealing a file, message, image, or video within another file, message, image, or video. Special programs are often used to hide messages in media files. If a hacker intercepts the message, all they see is the media. They don't know that there is a hidden message.

By combining a user's private encryption key and a hash of the data, a user can create a digital signature. A digital signature verifies that the data is legitimate and provides non-repudiation. This means that the sender cannot deny having sent the file.

q_cryp_concepts_hashing_secp8

Hashing is the process of converting one value into another using a mathematical algorithm like MD5 or SHA. This fixed length of data is called the hash.

Which of the following are true statements about hashing? (Select two.)

Answers:

- ***Hashing is used on data that does not need to be decrypted, such as a password.**
- ***A hash cannot be decrypted.**
- A hash is a string of bits that is randomly generated using a specific cipher.
- Hashing works by shifting each letter in the alphabet a certain number of spaces to the right or left.
- A hash verifies that the data is legitimate and provides non-repudiation.

Explanation:

The following are true statements about hashing:

- Hashing is used on data that does not need to be decrypted, such as a password.
- A hash cannot be decrypted.
- When a piece of data is run through a hashing algorithm, it always generates the same hash.

An encryption key is a string of bits that is randomly generated using a specific cipher, such as Advanced Encryption Standard (AES).

A Caesar cipher works be shifting each letter in the alphabet a certain number of spaces to the right or left.

A digital signature verifies that the data is legitimate and provides non-repudiation. This means that the sender cannot deny having sent the file.

q_cryp_concepts_stegano_secp8

Which of the following terms means a cryptography mechanism that hides secret communications within various forms of data?

Answers:

- Cryptanalysis
- ***Steganography**
- Algorithm
- Ciphertext

Explanation:

Steganography is the cryptography mechanism that hides secret communications within various forms of data.

Ciphertext is the encrypted form of a message that makes it unreadable to all but those the message is intended for.

Cryptanalysis is the method of recovering original data that has been encrypted without having access to the key used in the encryption process.

A cipher or algorithm is the process or formula used to convert a message or otherwise hide its meaning.

q_asys_sys_encrypt_asym_01_secp8

How many keys are used with asymmetric (public key) cryptography?

Answers:

- One
- ***Two**
- Three
- Four

Explanation:

Public key (asymmetric) cryptography uses two keys. One key is referred to as the public key and the other as the private key. This key pair overcomes the difficulties associated with the secure distribution of private keys. The communicating parties do not need to share secret information, as only the public keys are shared. Public keys are associated with users through authentication, usually through a mutually trusted directory, such as a certificate authority. The sender transmits a confidential message using only the recipient's public key. The message can only be decrypted with the associated private key possessed solely by the recipient. Public key cryptography not only provides encryption but is the basis for authentication technologies such as digital signatures.

q_asys_sys_encrypt_asym_02_secp8

Which of the following algorithms are used in asymmetric encryption? (Select two.)

Answers:

- ***RSA**
- Twofish
- AES
- ***Diffie-Hellman**
- Blowfish

Explanation:

RSA and Diffie-Hellman are asymmetric algorithms. RSA, one of the earliest encryption algorithms, can also be used for digital signatures. The Diffie-Hellman Protocol was created in 1976 but is still in use today in technologies such as SSL, SSH, and IPsec.

AES is the most the most commonly-used symmetric (not asymmetric) cipher.

Twofish is a symmetric encryption algorithm that uses a single key to both encrypt and decrypt data and information.

Blowfish is a variable-length, symmetric, 64-bit block cipher.

q_asys_sys_encrypt_asym_03_secp8

A systems administrator for a large organization dealing with massive amounts of data must ensure the security of data stored on disks and during transfers over the network.

Recently, the organization has encountered challenges with encrypting and decrypting large volumes of data.

Which encryption method would be inefficient for encrypting a large amount of data on a disk or transporting it over a network?

Answers:

- ***Asymmetric encryption**
- Symmetric encryption
- Hybrid encryption
- Obfuscation techniques

Explanation:

While asymmetric encryption offers advantages in certain areas, it becomes inefficient when dealing with large volumes of data on disks or during network transport due to its computational overhead.

Asymmetric encryption requires more processing power and time when compared to symmetric encryption, which is more efficient for bulk data encryption.

Hybrid encryption combines the benefits of both asymmetric and symmetric encryption but does not alleviate the inefficiencies of asymmetric encryption for large data encryption.

Obfuscation techniques are not encryption methods and would not address the performance challenges associated with encrypting large amounts of data. Obfuscation involves making a message or data difficult to find, such as hiding messages in files.

q_asys_sys_encrypt_ephemeral_keys_secp8

Which of the following statements about ephemeral keys is true?

Answers:

- Ephemeral keys are long-term keys used for multiple sessions.
- Ephemeral keys are used to provide integrity, but not confidentiality.
- Ephemeral keys are used in symmetric encryption only.
- ***Ephemeral keys are temporary and used for a single session only.**

Explanation:

Ephemeral keys are temporary keys used for a single session only. They are generated for each session and discarded after use, which provides an additional layer of security and perfect secrecy.

Ephemeral keys are not long-term keys; they are temporary and used for a single session only.

Ephemeral keys are used to provide both integrity and confidentiality. They ensure that the session's encryption key cannot be discovered even if the long-term private key is compromised.

Ephemeral keys can be used in both symmetric and asymmetric encryption. In asymmetric encryption, they are often used in combination with long-term keys to provide perfect forward secrecy.

q_asys_sys_encrypt_hash_secp8

A receiver wants to verify the integrity of a message received from a sender. A hashing value is contained within the digital signature of the sender.

Which of the following must the receiver use to access the hashing value and verify the integrity of the transmission?

Answers:

- Receiver's private key
- Receiver's public key
- Sender's private key
- ***Sender's public key**

Explanation:

Digital signatures are created using the sender's private key. Therefore, only the sender's public key can be used to verify and open any data encrypted with the sender's private key.

The recipient's private and public keys are not involved in this type of cryptography situation. Often, the hashing value of a message is protected by the sender's private key (their digital signature). The recipient must extract the original hashing value.

q_asys_sys_encrypt_hybrid_crypto_secp8

Which of the following statements about hybrid cryptosystems is true?

Answers:

- Hybrid cryptosystems use only symmetric encryption for both key exchange and data transmission.
- In hybrid cryptosystems, the symmetric key is used to encrypt the public key.
- Hybrid cryptosystems use the same key for both encryption and decryption processes.
- ***In hybrid cryptosystems, the public key is used to encrypt the symmetric key which is then used for data encryption.**

Explanation:

In hybrid cryptosystems, the public key is used to encrypt the symmetric key. This encrypted symmetric key is then sent to the receiver who uses their private key to decrypt it. The decrypted symmetric key is then used for data encryption and decryption.

Hybrid cryptosystems use asymmetric encryption for key exchange and symmetric encryption for data transmission.

In hybrid cryptosystems, the public key is used to encrypt the symmetric key, not the other way around.

While it's true that symmetric encryption uses the same key for encryption and decryption, hybrid cryptosystems also involve asymmetric encryption which uses different keys for encryption and decryption.

q_asys_sys_encrypt_key_01_secp8

Mary wants to send a message to Sam in such a way that only Sam can read it.

Which key should be used to encrypt the message?

Answers:

- ***Sam's public key**
- Sam's private key
- Mary's public key
- Mary's private key

Explanation:

Sam's public key should be used to encrypt the message. Only the corresponding private key, which only Sam has, can be used to decrypt the message.

Mary cannot use Sam's private key because only Sam has that key. Anything encrypted with the private key can be decrypted by anyone with the public key.

Encrypting the message using Mary's private key would mean that anyone could read the data using Mary's public key. Encrypting with Mary's public key would mean that only Mary would be able to decrypt it using her private key.

q_asys_sys_encrypt_key_02_secp8

Mary wants to send a message to Sam. She wants to digitally sign the message to prove that she sent it.

Which key would Mary use to create the digital signature?

Answers:

- Sam's public key
- Sam's private key
- Her public key
- ***Her private key**

Explanation:

Mary should use her private key to create the digital signature. This proves that only Mary could have sent the message because only Mary has access to her private key. Sam would use Mary's public key to verify the digital signature.

Use Sam's public key to encrypt a message that only Sam should be able to read. Only the corresponding private key, which only Sam has, can be used to decrypt the message.

Mary cannot use Sam's private key because only Sam has that key. Anything encrypted with the private key can be decrypted by anyone with the public key. Encrypting the message with Mary's public key would mean that only Mary would be able to decrypt it using her private key. But she could not prove where the message came from because anyone has access to Mary's public key.

q_asys_sys_encrypt_key_03_secp8

Which type(s) of key(s) are used in symmetric cryptography?

Answers:

- Two unique sets of key pairs
- A single key pair
- ***A shared key**
- A unique key for each participant

Explanation:

Symmetric cryptography uses a shared key. Both communication partners must be in possession of the same key in order to exchanged encrypted data.

Asymmetric cryptography uses a unique key pair for each participant. This key pair consists of a public key and a private key.

q_asys_sys_encrypt_key_04_secp8

How many keys are used with symmetric key cryptography?

Answers:

- ***One**
- Two
- Four
- Five

Explanation:

Private key, or symmetric, cryptography uses a single shared key. Both communicating parties must possess the shared key to encrypt and decrypt messages. The biggest challenge to symmetric cryptography is the constant need to protect the shared private key. This protection must be applied at all times, including during the initial transmission of the shared key between the parties.

q_asys_sys_encrypt_secret_01_secp8

Above all else, what must be protected to maintain the security and benefit of an asymmetric cryptographic solution, especially if it is widely used for digital certificates?

Answers:

- ***Private keys**
- Public keys
- Cryptographic algorithm
- Hash values

Explanation:

The strength of an asymmetric cryptographic system lies in the secrecy and security of its private keys. A digital certificate and a digital signature are little more than unique applications of a private key. If the private keys are compromised for a single user, for a secured network, or for a digital certificate authority, the entire realm of trust is destroyed.

The public key cannot be used to decrypt ciphertext (only the private key), and is less important that the private key.

Asymmetric encryption can be implemented using a number of cryptographic algorithms. However, no matter what algorithm is being used, without the private key the data cannot be decrypted.

Hash values are part of the process of using hashing algorithms. However, as with cryptographic algorithms, without the private key the encrypted data cannot be decrypted.

q_asys_sys_encrypt_secret_02_secp8

The success of asymmetric encryption is MOST dependent upon which of the following?

Answers:

- The integrity of the individuals who created the cryptosystem
- The secrecy of the algorithm
- The complexity of the ciphertext
- ***The secrecy of the key**

Explanation:

The strength of an asymmetric encryption system lies in the secrecy and security of its private keys.

While it's important to trust those who create a cryptosystem, if you are using a cryptosystem it is probably because it has already been recognized as a trusted system by the industry and other companies.

The strength of a cryptosystem should not be in the secrecy of the algorithm. This means that the algorithm is usually published and can be scrutinized for weaknesses.

No matter how complex the algorithm, if the private keys are not kept secure, the encrypted data is at risk.

q_asys_sys_encrypt_sym_01_secp8

Which of the following algorithms are used in symmetric encryption? (Select two.)

Answers:

- ***DES**
- RSA
- ***Blowfish**
- Diffie-Hellman
- ECC

Explanation:

DES and Blowfish are symmetric encryption algorithms.

RSA, Diffie-Hellman, and ECC are asymmetric encryption algorithms.

q_asys_sys_encrypt_sym_02_secp8

If a message sender encrypts a message with a key and a message receiver decrypts it using the same key, which type of key exchange is taking place?

Answers:

- ***Symmetric**
- Asymmetric
- Digital signature
- Counter mode

Explanation:

A symmetric key is when the sender uses a public key to encrypt a message and the recipient uses that same public key to decrypt it.

An asymmetric key is where the sender's and receiver's keys are both different for the encryption and decryption processes.

Using counter mode, both the sender and recipient access a reliable counter that computes a new shared value each time a ciphertext block is exchanged.

A digital signature is a mathematical scheme for demonstrating the authenticity of digital message or document.

q_asys_sys_encrypt_sym_03_secp8

A newly launched online store wants to secure transactions between the store and customers. The store must guarantee the authenticity of transactions, provide confidentiality, and ensure that only authorized recipients can access the purchase details.

Which cryptographic technique would best meet these requirements?

Answers:

- Symmetric encryption
- Hashing techniques
- ***Asymmetric encryption**
- Hybrid encryption

Explanation:

Asymmetric encryption uses a pair of keys - • public and private. The online store can encrypt the transaction details with the customer's public key, ensuring that only the customer, who holds the corresponding private key, can decrypt and access the details.

Symmetric encryption provides confidentiality fails to inherently provide authentication or non-repudiation, as the same key is for both encryption and decryption.

Hashing techniques verify the integrity of the data, but they do not provide encryption or the capability to validate the sender's authenticity.

While hybrid encryption provides the benefits of both symmetric and asymmetric encryption, it is more complex and allows for secure key exchange and efficient data encryption.

q_asys_sys_encrypt_weak_01_secp8

Which of the following encryption mechanisms offers the least security because of weak keys?

Answers:

- AES
- TwoFish
- IDEA
- ***DES**

Explanation:

DES offers the least encryption security of all the cryptography systems in this list. DES has a limitation of 56-bit keys, the weakest of those listed here. The strength of a cryptosystem lies not only in long keys but in the algorithm, initialization vector or method, the proper use of the keyspace, and the protection and management of keys.

AES (128-, 192-, and 256-bit keys), TwoFish (up to 256-bit keys), and IDEA (128-bit keys) all support stronger keys than DES.

q_asys_sys_encrypt_weak_02_secp8

Which of the following is the weakest symmetric encryption method?

Answers:

- ***DES**
- AES
- 3DES
- Twofish
- Blowfish

Explanation:

DES was one of the first symmetric encryption methods and is now obsolete (known weaknesses can be used to break the encryption).

3DES improves upon DES by applying the encryption three times. It is an acceptable alternative to DES. AES is stronger and faster than 3DES when implemented with a large key size (256 bits).

Blowfish and Twofish were alternatives to DES, but AES was chosen to replace DES.

q_cryp_algorithm_ctr_secp8

There are several block cipher modes of operation that can be utilized depending on the application or use.

Which of the following block cipher modes of operation uses a nonce combined with a counter that is encrypted?

Answers:

- ***Counter Mode (CTR)**
- Cipher Block Chaining (CBC)
- Cipher Feedback Mode (CFB)
- Electronic Code Book (ECB)

Explanation:

Instead of using an initialization vector, CTR uses a nonce combined with a counter that is encrypted. A nonce is a random string that is used for all blocks during the encryption process. The encrypted output of the nonce and counter is then XORed with the plaintext to create the ciphertext.

Cipher Block Chaining (CBC) uses an initialization vector (IV). The IV is a starting variable that is XORed with the plaintext of the current block to encrypt the data.

Cipher Feedback Mode (CFB) also uses an IV, but instead of using it on the plaintext, the IV is encrypted first. That output is then XORed with the plaintext to create the block of ciphertext.

Electronic Code Book (ECB) is the simplest mode of operation. Each block of plaintext data is encrypted separately. Blocks of data can be encrypted simultaneously allowing for faster encryption.

q_cryp_algorithm_gcm_secp8

Which of the following block cipher modes of operation provides both encryption and authentication?

Answers:

- ***Galois Counter Mode (GCM)**
- Cipher Block Chaining (CBC)
- Cipher Feedback Mode (CFB)
- Electronic Code Book (ECB)

Explanation:

The Galois Counter Mode (GCM) provides both encryption and authentication. All other block cipher modes of operation are unauthenticated forms of encryption.

Cipher Block Chaining (CBC) uses an initialization vector (IV). The IV is a starting variable that is XORed with the plaintext of the current block to encrypt the data.

Cipher Feedback Mode (CFB) also uses an IV, but instead of using it on the plaintext, the IV is encrypted first. That output is then XORed with the plaintext to create the block of ciphertext.

Electronic Code Book (ECB) is the simplest mode of operation. Each block of plaintext data is encrypted separately. Blocks of data can be encrypted simultaneously allowing for faster encryption.

q_cryp_algorithm_homomorphic_secp8

Which of the following types of encryption is specifically designed to allow data to be worked on without decrypting it first?

Answers:

- Lightweight cryptography
- ***Homomorphic encryption**
- Block cipher
- Stream cipher

Explanation:

An inherent risk with all encryption is that for the data to be worked on (computation on ciphertexts), it must first be decrypted. Homomorphic encryption addresses this concern by allowing data to be worked on without decrypting it first.

In 2018, NIST began the process to standardize encryption algorithms called lightweight cryptography. Lightweight cryptography is meant to be used on Internet of Things (IoT) devices.

Many symmetric encryption algorithms use the block cipher method. Instead of encrypting our data one bit at a time, a block cipher encrypts the data one chunk at a time.

A stream cipher is a symmetric encryption method that encrypts data one bit at a time. The stream cipher is based on the one-time pad (OTP) concept, which was used extensively during World War 2.

q_cryp_algorithm_lightweight_secp8

Which of the following types of encryption is specifically designed to be used on Internet of Things (IoT) devices?

Answers:

- ***Lightweight cryptography**
- Homomorphic encryption
- Block cipher
- Stream cipher

Explanation:

In 2018, NIST began the process to standardize encryption algorithms called lightweight cryptography. Lightweight cryptography is meant to be used on Internet of Things (IoT) devices.

An inherent risk with all encryption is that for the data to be worked on (computation on ciphertexts), it must first be decrypted. Homomorphic encryption addresses this concern by allowing data to be worked on without decrypting it first.

Many symmetric encryption algorithms use the block cipher method. Instead of encrypting our data one bit at a time, a block cipher encrypts the data one chunk at a time.

A stream cipher is a symmetric encryption method that encrypts data one bit at a time. The stream cipher is based on the one-time pad (OTP) concept, which was used extensively during World War 2.

q_cryp_algorithm_otp_secp8

Which of the following are true concerning the one-time pad (OTP) concept on which a streaming cipher is based? (Select two.)

Answers:

- ***OTP uses a symmetric encryption key that is the same length as the data being encrypted.**
- OTP encrypts the data one chunk at a time.
- ***OTP demonstrates what is called perfect secrecy.**
- OTP is meant to be used on Internet of Things (IoT) devices.
- OTP allows data to be worked on without decrypting it first.

Explanation:

The following is true concerning the OTP concept:

- OTP uses a symmetric encryption key that is the same length as the data being encrypted.
- OTP demonstrates what is called perfect secrecy.
- The encryption key is used to encrypt the data using a process called XORing.
- OTP only provides perfect secrecy if the secret key is only used once.
- OTP works well when dealing with small messages, but it becomes impractical when dealing with large data due to the keys being so extensive.

A block cipher encrypts data one chunk at a time.

Lightweight cryptography is meant to be used on Internet of Things (IoT) devices.

Homomorphic encryption allows data to be worked on without decrypting it first.

q_blockchain_business_processes_secp8

A company's leadership team has decided to integrate blockchain technology into its business processes to increase security, transparency, and efficiency. They have consulted a team of IT professionals to develop a secure communication infrastructure that supports their blockchain initiative.

When developing a robust and secure communication infrastructure for the implementation of blockchain technology, what factors should the IT professionals consider to maintain the integrity of the open public ledger and ensure the secure exchange of data across the network? (Select two.)

Answers:

- ***Implementing secure, encrypted communication protocols.**
- ***Ensuring the public ledger is accessible and transparent to all network members.**
- Restricting access to the public ledger to senior management only.
- Setting up data obfuscation techniques for the ledger entries.
- Ensuring that each transaction is securely stored on only one computer.

Explanation:

In a blockchain network, it is crucial to have

- Secure, encrypted communication protocols to maintain data integrity during transport.
- Transparency to keep the open public ledger accountable to all network members.

The public ledger in a blockchain system should be transparent to all members of the network, not only to the senior management.

While obfuscation can be a method to protect code, it does not typically get used within the entries of a blockchain's public ledger, which relies on transparency and validation across the network.

Every node in a blockchain contains a copy of the ledger. If a hacker changed a transaction, it wouldn't matter because a copy of the transaction is stored on many other computers.

q_blockchain_process_bitcoin_secp8

Which of the following was the first big use of blockchain cryptography process?

Answers:

- ***Cryptocurrency**
- Contracts
- Records
- Food

Explanation:

While the blockchain cryptography method can be used for items such as tracking contracts, records, and food products, the first big use of the blockchain was the Bitcoin cryptocurrency.

q_blockchain_process_steps_secp8

Blockchain is a unique and increasingly popular implementation of cryptography. A blockchain is a decentralized and distributed ledger that records and verifies transactions between two parties.

The list on the left describes each step a block goes through as part of the blockchain cryptographic process.

From the list on the left, drag a description to its proper step order on the right.

Answers:

- User1 requests a transaction with User2.
- The block is added to the chain.
- The network users verify the transaction is valid.
- The transaction is represented online as a block.
- The contents of the transaction move to User2.
- The block is distributed to everyone on a peer-to-peer network.

Explanation:

The following (in order) are the steps each block goes through as part of the blockchain cryptographic process:

1. User1 requests a transaction with User2. The request is made using User1's personal secret key and User2's public key.
2. The transaction is represented online as a block.
3. The block is distributed to everyone on a peer-to-peer network.
4. The network users verify the transaction is valid.
5. The block is added to the chain. This provides a indisputable and transparent record of the transaction.
6. The contents of the transaction move to User2.

q_cryp_attacks_collision_secp8

An incident response team detected a threat actor attempting to find two different inputs that produce the same hash value in the company's secure communication system.

Which of the following BEST describes this type of attack?

Answers:

- ***Collision attack**

- Birthday attack
- Brute-force attack
- Side-channel attack

Explanation:

In a collision attack, an attacker tries to find two different inputs that produce the same hash value.

Birthday attacks focus on the probability that two users will have the same hashed value, as opposed to deliberately trying to find two different inputs producing the same hash.

Brute-force attacks can apply in cryptographic systems but do not involve finding two different inputs that produce the same hash value.

Side-channel attacks are techniques used to extract secret cryptographic keys from a system by analyzing physical information leakage, such as time to execute cryptographic algorithms, power consumption, or even sound. They do not necessarily pertain to creating hash collisions.

q_cryp_attacks_dictionary_secp8

Which type of password attack employs a list of pre-defined passwords that it tries against a login prompt?

Answers:

- Collision attack
- Birthday attack
- ***Dictionary attack**
- Downgrade attack

Explanation:

A dictionary attack is a type of brute-force attack. A hacker uses a list of words and phrases to try to guess the decryption key.

- Dictionary attacks work well if weak passwords are used.
- Using longer and uncommon passphrases is the best way to secure data against these attacks.

A collision attack tries to find two inputs that produce the same hash value. This type of attack is often used on digital signatures.

A birthday attack combines a collision attack and brute-force attack. The name is taken from the birthday probability math problem.

A downgrade attack forces the system to use an older, less secure communication protocol.

q_cryp_attacks_downgrade_secp8

Which of the following cryptographic attacks uses SSL exploitation as a common implementation of this attack?

Answers:

- ***Downgrade attack**
- Birthday attack

- Collision attack
- Dictionary attack

Explanation:

A downgrade attack forces the system to use an older, less secure communication protocol. SSL exploitation is a common implementation of this type of attack.

A birthday attack combines a collision attack and brute-force attack. The name is taken from the birthday probability math problem.

A collision attack tries to find two inputs that produce the same hash value. This type of attack is often used on digital signatures.

A dictionary attack is a type of brute-force attack. The hacker uses a list of words and phrases to try to guess the decryption key.

q_cryp_attacks_quantum_computing_secp8

What is the key difference between classical computing and quantum computing that makes quantum computing exponentially more powerful?

Answers:

- Quantum computers use electricity, while classical computers do not.
- Quantum computers use bits, while classical computers use qubits.
- ***Quantum computers can process multiple possibilities simultaneously, while classical computers cannot.**
- Quantum computers are smaller than classical computers.

Explanation:

Quantum computers can process multiple possibilities simultaneously, while classical computers cannot is correct. Quantum computers use qubits, which can exist as both a 1 and 0 at the same time. This allows quantum computers to process multiple possibilities simultaneously, making them exponentially more powerful than classical computers.

Both quantum and classical computers use electricity. The difference lies in the way they process information.

Classical computers use bits (binary digits) that can be either a 0 or a 1. Quantum computers use qubits, which can exist in multiple states at once.

The size of a computer does not determine its computational power. Quantum computers can be larger, smaller, or the same size as classical computers.

3.2 Cryptography Implementations

As you study this section, answer the following questions:

- How can cryptography support the goals of information security?
- What cryptographic information is stored in a digital certificate?
- Why is reusing encryption keys considered a weakness?
- What are the potential consequences if a company loses control of a private key?

- What mechanism informs clients about suspended or revoked keys?
- What functionality does a Trusted Platform Module (TPM) chip provide?

In this section, you will learn to:

- Verify a device for TPM.

The key terms for this section include:

Term	Definition
Obfuscation	A technique that essentially hides or camouflages code or other information so that it is harder to read by unauthorized users.
Steganography	A technique for obscuring the presence of a message, often by embedding information within a file or other entity.
Data masking	A de-identification method where generic or placeholder labels are substituted for real data while preserving the structure or format of the original data.
Tokenization	A de-identification method where a unique token is substituted for real data.
Key management system	In public key infrastructure (PKI), procedures and tools that centralizes generation and storage of cryptographic keys.
Trusted Platform Module (TPM)	Specification for secure hardware-based storage of encryption keys, hashed passwords, and other user- and platform-identification information.
Application programming interface (API)	Methods exposed by a script or program that allow other scripts or programs to use it. For example, an API enables software developers to access functions of the TCP/IP network stack under a particular operating system.
Secure enclave	CPU extensions that protect data stored in system memory so that an untrusted process cannot read it.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions. • Encryption ○ Level

	▪ Full-disk ○ Asymmetric ○ Symmetric ○ Key exchange • Tools ○ Trusted Platform Module (TPM) ○ Hardware security module (HSM) ○ Key management system ○ Secure enclave • Obfuscation ○ Steganography ○ Hashing • Digital signatures
TestOut Security Pro	4.0 Data Security 4.2 Implement Encryption Technologies 4.2.1 Encrypt data communications 4.2.2 Encrypt files

3.2.1 Combining Cryptographic Methods (Lesson Video)

Transcript:

Using cryptography lets you secure your data against unauthorized users, but in some cases, it's beneficial if you combine different cryptographic methods to increase data security.

In this video, we'll look at hybrid models, digital signatures, and combining encryption with steganography.
Hybrid cryptographic models combine the use of symmetric and asymmetric keys to encrypt data.
Asymmetric encryption methods provide an easier way to exchange keys safely, but are not very efficient because they rely on some complicated mathematical computations. Symmetric encryptions are efficient, but because it can be a pain to get the keys out, they're not very convenient. A hybrid cryptosystem combines the convenience of asymmetric methods and the efficiency of symmetric methods.
One example of this combination is the Microsoft Encrypting File System, or EFS. EFS was introduced in version 3.0 of NTFS. It has been included in every version of Windows except Home edition since Windows 2000.
EFS works by encrypting a file with a special symmetric key called the File Encryption Key, or FEK. The user's public asymmetric key is then used to encrypt the FEK.
During decryption, the user's private key is used to decrypt the FEK, which is then used to decrypt the file.
By combining the security of the symmetric keys and the convenience of the asymmetric keys, Microsoft has made it easy for users to encrypt their data. The main drawback of EFS is that the user's private key is essentially their password. This means that the encryption is only as strong as the user's password.
One very common practice that combines cryptographic methods
is the digital signature. A digital signature combines the hash of a file and a user's private key to electronically sign a document, providing an authentic and non-repudiation file.
For example, if we had a large file that needed to be digitally signed, we could use asymmetric encryption to simply encrypt it. However, that process is slow and could generate a very large file size. Instead, we can create a hash of the file.
The three main hashing algorithms used today are SHA-1, which generates a 128-bit key, MD5, which generates a 160-bit key, and SHA-2, which is also commonly referred to as SHA-256. It creates a 256-bit key.
Whichever algorithm is used, the hash generates a fingerprint of the file. Now, instead of encrypting the file itself, we can encrypt the hash that was generated using our private key. We combine that encrypted hash with the file to digitally sign it.

Now, when we send the message to the recipient, she'll first generate a hash of the file. She'll then use our public key to decrypt the hash that we generated and compare it with theirs. If the hashes match, then they can be assured that the message is legitimate and hasn't been altered.
We can take this one step further and encrypt the file itself along with the digital signature using a digital envelope. To do this, we would use the recipient's public key to encrypt the data. The ciphertext, hash, and digital signature are all combined and sent together. The recipient uses our public key to decrypt the hash to authenticate the message. Then they use their private key to decrypt the ciphertext.
By combining the digital signature with asymmetric encryption, we get authentication, confidentiality, integrity, and non-repudiation.
An interesting combination of cryptographic methods is using asymmetric encryption with steganography. There's been a fairly large amount of recent research on ways to combine these methods effectively.
Let's look at how this would work.
We start by encrypting our plaintext with our private key, generating our ciphertext. Next, we hide the ciphertext inside of a media file, such as an image.
When our recipient receives the file, she'll first extract the ciphertext and then use the public key to decrypt the ciphertext.
By combining these two methods, we increase the security of a message. Anyone intercepting the message would need to know that there's an encrypted file inside the image file before attempting to crack the encryption.
And that's it for this lesson. In this video, we discussed the different ways of combining cryptographic methods. We looked at how hybrid cryptographic models combine the efficiency of symmetric methods and the convenience of asymmetric methods and how EFS takes advantage of this combination. Then we discussed digital signatures and how they can be used to provide message authentication and non-repudiation. Finally, we went over combining encryption with steganography and how this combination can be used to hide ciphertext inside of a media file.

3.2.2 Hardware-Based Encryption Devices (Lesson Video)

Transcript:

Many encryption systems are managed by the operating system or other software, but there are also some hardware options for encryption. In this lesson, we'll look at the two of the most common hardware options, the TPM and HSM.

The Trusted Platform Module, or TPM, is a physical chip that resides on the motherboard. It's responsible for providing some cryptographic services. Using a hardware chip means that the encryption system itself can't be attacked by malicious software. The latest version, TPM 2.0, was released in 2014. TPM chips can be used for checking key system components at startup and for protecting and generating encryption keys and passwords.
One key function of the TPM chip is to check key system components during startup. When the system is turned on, the TPM checks to make sure everything seems normal. If the TPM detects anything out of sorts, it doesn't allow the system to boot, which prevents data from being extracted.
The TPM is also used to store and generate encryption keys. In essence, TPM provides an encrypted lockbox where user passwords, encryption keys, and digital certificates can be kept safe.
For example, Windows 10 can pull these keys directly from the TPM without loading them into the RAM, where they would be more vulnerable to attack.
When encryption keys are generated, they usually need a random number to be generated too. This is a weak spot in encryption since most software number generators contain patterns and are not truly random.
The TPM chip can be used to generate completely random numbers for the encryption process, which increases the effectiveness of the encryption keys.
A TPM chip can also greatly increase the security of a system. Some popular Windows features such as BitLocker and Credential Guard rely on the TPM chip to perform at their most secure level.
A hardware security module, or HSM, is very similar to a TPM. You might also hear these devices referred to as a Personal Computer Security Module or Secure Application Module. One of key differences is that an HSM is a removable device, whereas the TPM is built into the motherboard.
One major benefit of HSMs is that they're more powerful than TPM chips. HSMs are designed to be powerful cryptographic devices with better hardware capabilities. This allows HSMs to perform multiple security operations for multiple users. Some HSM devices can even be attached to a server to handle cryptographic functions for an entire network.

HSM devices perform many of the same functions as a TPM, such as generating and storing encryption keys, but they can also be used to generate and validate digital signatures and generate keys used in smart cards. Consumer-grade HSM devices have become much more affordable over the past few years, making them more accessible than before.
And that's it for this lesson. In this video, we've covered two options for hardware-based encryption. The TPM chip is built into the motherboard. It checks key hardware components during startup and generates and stores encryption keys. And HSMs are add-on devices that perform many of the same functions as a TPM chip, but are much more powerful.

3.2.3 Verify Device for TPM (Demo Video)

Transcript:

TPM, or Trusted Platform Module is a security device used to secure hardware. One of the most common ways is using BitLocker with a TPM. When combined, BitLocker encryption and a TPM provide more security than standard software encryption alone. To determine if your device has a TPM, you can access the BIOS. Look for an option related to Trusted Computing. Within this section, you'll find an option to enable or disable TPM. Some BIOS settings might not explicitly mention TPM, so Trusted Computing could be the option you're seeking if you don't see TPM mentioned. This concept should apply to most BIOS settings; you'll just need to locate the menu based on your motherboard model.

With newer versions of Windows, such as Windows 11, TPM 2.0 is required. Most computers manufactured within the last 5 years should have this capability. If your computer is older, there's a chance it may have an older TPM module or no TPM at all. In Windows, there are several ways to verify TPM. You can navigate to the Device Manager and go to Security Devices, where you'll find a listing for TPM. If Windows cannot detect a TPM, you'll see nothing in this section. You can also check the status of this device to ensure it's functioning correctly. The key thing to look for is to ensure there are no question marks or errors associated with this device. If there are, it may indicate a device or driver issue that needs attention.
Another method to check your TPM is by examining your security processor in Windows Security. Access Windows Security by typing security in the search bar, then navigate to Device Security and click on Security Processor Details. If your TPM is functioning properly, you will see related details here. In case you encounter issues with your TPM, you have the option to troubleshoot and clear TPM information. Note that you should only clear TPM if you intend to remove previously configured TPM data. There could be situations where a computer was previously used with BitLocker, and clearing the TPM may become necessary.
Since we're discussing BitLocker, let's explore some settings related to it. If you navigate to BitLocker from the Start menu, you will notice that encryption is not currently enabled. One area not accessible via a Windows search bar is TPM management. If you are familiar with older versions of Windows, this window should look familiar. It offers similar options to the troubleshooting area in security, such as clearing your TPM or preparing it for use. Since the status is set to Ready to Use, you should be able to enable BitLocker and utilize TPM for enhanced security.
That concludes this demonstration. In this demo, we have shown you various methods for verifying the presence of a TPM and ensuring it is ready for use with BitLocker encryption.

3.2.4 Cryptographic Implementation Facts

This lesson covers the following topics:

- Hybrid models
- Obfuscation
- Encryption with steganography
- Key management
- Secure enclaves
- Cryptoprocessors

Hybrid Models

Operating systems, applications, and other components of information systems typically use a hybrid cryptography system. A hybrid cryptography system combines the strengths of hashing, symmetric encryption, and asymmetric encryption, depending on the needs of the project or service. An example of these strengths are:

- Use symmetric encryption for fast and efficient encryption of bulk data.
- Use hashing to verify message integrity.
- Use asymmetric encryption for authentication and non-repudiation.
- Use asymmetric encryption for secure exchange of symmetric encryption keys (for example, by encrypting the key used for symmetric encryption prior to sharing the key with the recipient). Using asymmetric cryptography for encryption is best for small pieces of data.

A hybrid cryptosystem combines the efficiency of symmetric methods and the convenience of asymmetric methods. One example of this is Microsoft's Encrypting File System, or EFS. Microsoft's EFS uses the following steps to encrypt data:

1. A file is encrypted using a File Encryption Key (FEK).
2. The FEK is encrypted with the user's public asymmetric key.
3. The file is sent to the intended recipient.
4. The user's private key is used to decrypt the FEK.
5. The FEK is used to decrypt the file.

One of the biggest weaknesses of the EFS is that the user's private key is essentially their user password. If the password is weak, the encryption will also be weak.

Obfuscation

Obfuscation is the art of making a message or data difficult to find. It is security by obscurity, which is normally deprecated. There are some uses for obfuscation technologies, however:

- Steganography (literally meaning "hidden writing") embeds information within an unexpected source, a message hidden in a picture, for instance. The container document or file is called the *covertext* . The message can be encrypted by some mechanism before embedding it, providing confidentiality. The technology can also provide integrity or non-repudiation; for example, it could show that something was printed on a particular device at a particular time, which could demonstrate that it was genuine or fake, depending on the context.
- Data masking can mean that all or part of the contents of a database field are redacted by substituting all character strings with "x," for example. A field might be partially redacted to preserve metadata for analysis purposes. For example, in a telephone number, the dialing prefix might be retained, but the subscriber number is redacted. Data masking can also use techniques to preserve the original format of the field.
- Tokenization means that all or part of the value of a database field is replaced with a randomly generated token. The token is stored with the original value on a token server or token vault, separate from the production database. An authorized query or app can retrieve the original value from the vault if necessary, so tokenization is reversible. Tokenization is used as a substitute for encryption because, from a regulatory perspective, an encrypted field has the same value as the original data.

Data masking and tokenization are used for de-identification. De-identification obfuscates personal data from databases so that it can be shared without compromising privacy.

Encryption with Steganography

There has been an increased amount of research done on how to best combine encryption with steganography. The process essentially follows the steps below:

1. Encrypt plaintext with a private key to generate ciphertext.

2. The ciphertext is hidden inside a media file, such as an image, using steganography.
3. The recipient extracts the ciphertext and decrypts it using the matching public key.
4. Because the ciphertext is hidden in the image file, someone intercepting the message would have to know it is there before being able to decrypt it.

Key Management

Key management refers to operational considerations for the various stages in a key's lifecycle. A key's lifecycle may involve the following stages:

- **Key Generation** — creates an asymmetric key pair or symmetric secret key of the required strength using the chosen cipher.
- **Storage** — prevents unauthorized access to a private or secret key and protects against loss or damage.
- **Revocation** — prevents use of the key if it is compromised. If a key is revoked, any data that was encrypted using it should be re-encrypted using a new key.
- **Expiration and Renewal** — gives the certificate a "shelf-life," increasing security. Every certificate expires after a certain period. Certificates can be renewed with the same key pair or with a new key pair.

A decentralized key management model means that keys are generated and managed directly on the computer or user account that will use the certificate. This does not require any special setup, so it is easy to deploy. However, it makes the detection of key compromises more difficult.

Some organizations prefer to centralize key generation and storage using a tool such as a key management system . In one type of cryptographic key management system, a dedicated server or appliance is used to generate and store keys. When a device or app needs to perform a cryptographic operation, it uses the Key Management Interoperability Protocol (KMIP) to communicate with the server.

Secure Enclaves

A key pair or secret key can be generated and stored in the file system on a desktop or server computer running a general-purpose OS. This has a number of drawbacks, however:

- **A cryptographic key needs to be generated using a random process** . A key generation system with a high degree of disorder—or entropy—ensures that any value from the possible keyspace has the same chance of being selected as any other. Unfortunately, computer hardware and software are extremely low entropy—computers process instructions in an entirely deterministic way. A computer can use pseudo RNG (PRNG) software that is still deterministic but able to approximate a high level of disorder. Better security is obtained by true random number generator (TRNG) hardware. This uses a source of entropy, such as noise or air movement, as a nondeterministic seed for generating the key value.
- **A key stored in the file system is only as secure as any other file.** It could easily be compromised via the user credential or physical theft of the device. It is also difficult to ensure that key access is fully audited. Ideally, cryptographic storage is tamper-evident. This means that it is known immediately when a private or secret key has been compromised, and it can be revoked and any ciphertexts re-encrypted with a new key.

```
GnuPG needs to construct a user ID to identify your key.

Real name: James Pengelly
Email address: jpengelly@comptia.org
Comment:
You selected this USER-ID:
    "James Pengelly <jpengelly@comptia.org>"

Change (N)ame, (C)omment, (E)mail or (O)kay/(Q)uit? o
We need to generate a lot of random bytes. It is a good idea to perform
some other action (type on the keyboard, move the mouse, utilize the
disks) during the prime generation; this gives the random number
generator a better chance to gain enough entropy.
We need to generate a lot of random bytes. It is a good idea to perform
some other action (type on the keyboard, move the mouse, utilize the
disks) during the prime generation; this gives the random number
generator a better chance to gain enough entropy.
gpg: /home/kali/.gnupg/trustdb.gpg: trustdb created
gpg: directory '/home/kali/.gnupg/openpgp-revocs.d' created
gpg: revocation certificate stored as '/home/kali/.gnupg/openpgp-revocs.d/C3E
1D43D17CBCC7C80A3D4889564BC94BD4E1D99.rev'
public and secret key created and signed.

pub   rsa2048 2023-05-31 [SC] [expires: 2025-05-30]
      C3E1D43D17CBCC7C80A3D4889564BC94BD4E1D99
uid                      James Pengelly <jpengelly@comptia.org>
sub   rsa2048 2023-05-31 [E] [expires: 2025-05-30]
```

Pseudo RNG working during key generation using GPG. This method gains entropy from user mouse and keyboard usage.

These drawbacks can be addressed using a cryptoprocessor for key generation and storage.

Cryptoprocessors

A cryptoprocessor performs operations such as decryption and signing on behalf of apps. This means that the key material never leaves the cryptoprocessor. Because it is dedicated to a single function, the cryptoprocessor hardware has a smaller attack surface than a general-purpose computer.

Hardware security module appliance in a rack-mounted form factor. (Images © 123RF.com.)

There are two main ways of implementing cryptoprocessor hardware:

Cryptoprocessor Hardware	Description

Trusted Platform Module (TPM)	A TPM is a cryptoprocessor implemented as a module within the CPU on a computer or mobile device. TPMs are produced to different version specifications, with versions 1.2 and 2.0 in current use. Version 2.0 is not backward compatible with version 1.2. Beginning with Windows 10 version 1607, Microsoft required that TPM 2.0 be enabled by default on all new computers.
	• A TPM is required to check the integrity of startup files and components in BitLocker implementations.
	  ○ The TPM generates a hash of the startup files to verify the integrity of those files.
	  ○ Additionally, the TPM creates a hash of system components. This hash acts as a validation check of the system to ensure that system components have not changed. The hash can also be used to uniquely identify the system.
	• Windows Credential Guard requires the computer to have a TPM chip installed.
	  ○ A TPM provides protection for virtual-based security encryption keys that are stored in the firmware. This helps protect against attacks involving a physically present user with BIOS access.
	• A TPM can generate truly random numbers, thus preventing entropy.
	• TPM provides full support for asymmetric encryption; therefore, it can generate public and private keys.
	• A TPM also provides encrypted storage for user passwords, encryption keys, and digital certificates.
	• Windows 10 can pull stored keys directly from the TPM without loading them into the RAM, where they would be more vulnerable to an attack.

Hardware security module (HSM)	An HSM is cryptoprocessor hardware implemented in a removable or dedicated form factor, including rack-mounted appliances, plug-in PCIe adapter cards, and USB-connected security keys. It is also possible to provision an HSM as a virtual appliance.

Cryptoprocessor Hardware	Description
	HSMs provide cryptographic functions such as: • Generate and store encryption keys. • Generate and validate digital signatures. • Generate keys used in smart cards. HSMs traditionally come in the form of a plug-in card or an external security device that can be attached directly to the computer system. These devices offer some benefits over TPM chips. • HSMs are more powerful and can perform more powerful cryptographic functions quicker. • HSMs can perform multiple cryptographic functions simultaneously. • HSMs can be attached to a network and handle cryptographic functions for multiple users across the network. Hardware Security Modules are also known as: • Personal Computer Security Module (PCSM) • Secure Application Module (SAM) • Hardware cryptographic devices • Cryptographic modules

Vendors can certify their products against the Federal Information Processing Standard 140 Level 2 (FIPS 140-2) to establish trust in the market.

Using a cryptoprocessor means that keys are not directly accessible via the file system. The cryptoprocessor interacts with applications that need to access the key via an application programming interface (API) that implements PKCS#11.

One vulnerability in this system is that decrypted data needs to be loaded into the computer's system memory (RAM) for applications to access it. This raises the potential for a malicious process to gain access to the data via some type of exploit. This vulnerability can be mitigated by implementing a secure enclave . A trusted execution environment (TEE) secure enclave, such as Intel Software Guard Extensions, can protect data stored in system memory so that an untrusted process cannot read it. A secure enclave is designed so that even processes with root or system privileges cannot access it without authorization. The enclave is locked to a list of one or more digitally signed processes.

3.2.5 Practice Questions (Section Quiz)

q_comb_cryp_encrypt_secp8

Which form of cryptography is BEST suited for bulk encryption because it is so fast?

Answers:

- Hashing cryptography
- Public key cryptography

- ***Symmetric key cryptography**
- Asymmetric cryptography

Explanation:

Symmetric cryptography is best suited for bulk encryption because it is much faster than asymmetric cryptography.

Hashing is not used for encryption; it is only used to verify the integrity of data.

Public key cryptography, also known as asymmetric cryptography, is best suited for small amounts of data. Often, asymmetric cryptography is used to exchange symmetric cryptography keys, and then the symmetric cryptography keys are used to encrypt communication traffic.

q_comb_cryp_expiration_date_secp8

You are a cybersecurity manager at a financial institution. Your team is responsible for managing the cryptographic keys used for secure transactions.

Recently, there has been an increase in attempted cyber attacks on your institution.

Which of the following key management strategies would be MOST effective in maintaining the security of your cryptographic keys under these circumstances?

Answers:

- You decide to generate new keys for each transaction, but keep the old keys stored in the system for future reference.
- You decide to revoke all current keys and generate new ones, informing all users to update their keys immediately.
- ***You decide to set an expiration date for all current keys and inform users that they will need to renew their keys after this date.**
- You decide to centralize key generation and storage, moving all keys to a single server for easier management.

Explanation:

Setting an expiration date for all current keys and requiring users to renew their keys after this date is the correct and proactive approach to key management in this scenario. This ensures that keys are regularly updated, reducing the chance that an attacker can use an old key to gain unauthorized access. It also provides a structured and predictable process for users, which can help prevent confusion or errors.

While generating new keys for each transaction can enhance security, keeping old keys stored in the system can pose a risk. If an attacker gains access to these old keys, they might be able to decrypt past transactions and gain access to sensitive information.

Revoking all current keys and generating new ones can be a drastic measure that may cause unnecessary disruption. It can also confuse users and lead to potential errors or security risks if not managed properly.

While centralizing key generation and storage can simplify management, it can also create a single point of failure. If the server where the keys are stored is compromised, all keys could be at risk. It's generally better to distribute keys across multiple secure locations.

q_comb_cryp_hsm_secp8

A cyber technician reduces a computer's attack surface by installing a cryptoprocessor that a plug-in PCIe adaptor card can remove.

What type of cryptoprocessor can support this requirement?

Answers:

- ***HSM**
- CRL
- PKI
- TPM

Explanation:

A hardware security module (HSM) meets the analyst's needs in this scenario. An HSM is a cryptoproccessor that implements hardware through a removable or dedicated form factor, such as plug-in peripheral component interconnect express (PCIe) adaptor cards.

Not viable in this scenario, a certificate revocation list (CRL) provides a summation of all revoked and suspended certificates and must be accessible to anyone relying on the validity of the certificate authority's certificates.

Public key infrastructure (PKI) is the framework that establishes trust in using public key cryptography to sign and encrypt messages via digital signatures.

A trusted platform module (TPM) is a cryptoprocessor implemented as a module within the CPU on a computer or mobile device.

q_comb_cryp_obfuscation_secp8

You are a cybersecurity analyst at a large corporation. Your team has been tasked with securing sensitive data within the company's database. One of the strategies you are considering is obfuscation.

Which of the following scenarios would be the most appropriate application of obfuscation?

Answers:

- You use obfuscation to hide the company's financial data within an image file on the company's public website.
- You use obfuscation to hide the source code of the company's proprietary software within a text document.
- ***You use obfuscation to hide employee personal data within a database field by substituting character strings with x.**
- You use obfuscation to hide the company's network architecture within a PDF document.

Explanation:

You use obfuscation to hide employee personal data within a database field by substituting character strings with *x* is correct. This is a form of obfuscation known as data masking, which is often used to protect sensitive data, such as personal information. By replacing actual data with placeholder characters, you can prevent unauthorized access to the data while maintaining the original data structure for authorized use.

You use obfuscation to hide the company's financial data within an image file on the company's public website.
is incorrect. While this is a form of obfuscation (specifically steganography), it is not a good practice to hide sensitive data like

financial information on a public website. Even though it's hidden within an image file, if someone knows what to look for, they could potentially extract and exploit this information.

Obfuscating source code can make it more difficult for unauthorized users to understand and misuse it, but it's not a foolproof method of protection. If the obfuscated code is reverse-engineered, the original code could be exposed. Additionally, obfuscation does not prevent someone from copying the obfuscated code and using it elsewhere.

While obfuscation could be used to hide details about the company's network architecture, it's not the best method for protecting this type of information. A more effective approach would be to implement robust network security measures, such as firewalls, intrusion detection systems, and regular network monitoring.

q_comb_cryp_secure_enclaves_01_secp8

You are a cybersecurity architect at a tech company that is developing a new mobile payment application. The application will handle sensitive user data including credit card information and personal identification numbers (PINs).

Which of the following strategies would best leverage the concept of secure enclaves to protect this sensitive data?

Answers:

- You decide to store all sensitive data in a secure enclave within the application, accessible only through a secure API.
- You decide to store all sensitive data in a secure enclave on the company's main server, accessible only by senior IT staff.
- ***You decide to store all sensitive data in a secure enclave on each user's device, accessible only with the user's unique PIN.**
- You decide to store all sensitive data in a secure enclave within the application, accessible to all application users.

Explanation:

Storing sensitive data in a secure enclave on each user's device and making it accessible only with the user's unique PIN is the best strategy. This approach leverages the concept of secure enclaves by isolating sensitive data on a per-user basis, and it adds an additional layer of security by requiring the user's unique PIN for access.

While storing sensitive data in a secure enclave within the application and making it accessible only through a secure API is a good practice, it does not fully leverage the concept of secure enclaves. The data is still potentially vulnerable if the application itself is compromised.

Storing all sensitive data in a secure enclave on the company's main server could create a single point of failure. If the server is compromised, all data could be at risk. Additionally, this does not leverage the concept of secure enclaves as a method of isolating sensitive data on a per-user basis.

Storing sensitive data in a secure enclave within the application and making it accessible to all application users does not provide adequate protection for the data. This approach does not leverage the concept of secure enclaves, which are designed to isolate sensitive data and limit access to it.

q_comb_cryp_secure_enclaves_02_secp8

As a cybersecurity expert, you are tasked with implementing a secure enclave in your company's new mobile banking application.

Which of the following statements best describes the primary function and benefit of a secure enclave in this context?

Answers:

- A secure enclave is a protected area within the application's code that prevents users from making unauthorized transactions.
- ***A secure enclave is a separate, isolated environment within the device's processor where sensitive data can be securely stored and processed.**
- A secure enclave is a network security tool that monitors and controls incoming and outgoing network traffic based on predetermined security rules.
- A secure enclave is a cloud-based storage system where encrypted data is stored and can only be accessed with the correct decryption key.

Explanation:

A secure enclave is a separate, isolated environment within the device's processor. It is designed to securely store and process sensitive data, even in the event that the rest of the device is compromised.

While a secure enclave does provide a level of protection, it is not located within the application's code and it does not directly prevent users from making unauthorized transactions. Its primary function is to securely store and process sensitive data.

Monitoring and controlling incoming and outgoing network traffic refers to a firewall, not a secure enclave. A firewall is a network security device that monitors and controls network traffic, while a secure enclave is a hardware-based security feature of a processor.

While a secure enclave does involve the storage of sensitive data, it is not a cloud-based storage system. A secure enclave is a hardware-based feature that exists within a device's processor.

q_comb_cryp_stenography_secp8

Combining encryption with steganography involves several steps.

From the list on the left, drag a description of a step or result in this process to the correct order on the right.

Answers:

- The ciphertext is hidden inside of a media file, such as an image, using steganography.
- Anyone intercepting the message would have to know its there before being able to decrypt it.
- Encrypt plaintext with a private key to generate ciphertext.
- The recipient extracts the ciphertext and decrypts it using the matching public key.

Explanation:

There has been an increased amount of research done on how to best combine encryption with steganography. The process essentially follows the steps below:

1. Encrypt plaintext with a private key to generate ciphertext.
2. The ciphertext is hidden inside of a media file, such as an image, using steganography.
3. The recipient extracts the ciphertext and decrypts it using the matching public key.
4. Because the ciphertext is hidden in the image file, someone intercepting the message would have to know its there before being able to decrypt it.

q_comb_cryp_tpm_01_secp8

What is the main function of a TPM hardware chip?

Answers:

- ***Generate and store cryptographic keys**
- Perform bulk encryption in a hardware processor
- Control access to removable media
- Provide authentication credentials on a hardware device

Explanation:

A Trusted Platform Module (TPM) is a hardware cryptoprocessor that resides on the motherboard. This hardware is used to store and generate cryptographic keys. These keys are used for encryption and authentication, but the TPM does not perform the actual encryption.

A smart card is a hardware device containing a digital certificate. The smart card can be used for authentication.

Special hardware processors perform bulk encryption in hardware rather than software. These processors typically encrypt data using AES or encrypt network traffic using IPsec.

q_comb_cryp_tpm_02_secp8

Which of the following functions are performed by a TPM?

Answers:

- ***Create a hash of system components**
- Perform bulk encryption
- Encrypt network data using IPsec
- Provide authentication credentials

Explanation:

A Trusted Platform Module (TPM) is a hardware cryptoprocessor that resides on the motherboard. This hardware is used to store and generate cryptographic keys. The TPM also generates hash values of system components. The hash value verifies that startup components have not been modified. Because each system has a unique hash value, the hash can also be used as a form of identification for the system.

Keys generated by the TPM are used for encryption and authentication, but the TPM does not perform the actual encryption.

A smart card is a hardware device containing a digital certificate. The smart card can be used for authentication.

Special hardware processors perform bulk encryption in hardware rather than software. These processors typically encrypt data using AES or encrypt network traffic using IPsec.

q_comb_cryp_tpm_03_secp8

A cyber security analyst wants to reduce the attack surface for a computer that contains top secret data. The analyst installs a cryptoprocessor as a module within the central processing unit (CPU) on the designated computer to accomplish this.

What type of cryptoprocessor is the analyst installing?

Answers:

- ***TPM**
- HSM
- CRLs
- PKI

Explanation:

A trusted platform module (TPM) meets the analyst's needs in this scenario. A TPM is a cryptoprocessor implemented as a module within the CPU on a computer or mobile device.

Not viable in this scenario, a hardware security module (HSM) is a cryptoprocessor that implements hardware through a removable or dedicated form factor, such as plug-in peripheral component interconnect express (PCIe) adaptor cards.

A certificate revocation list (CRL) provides a summation of all revoked and suspended certificates and must be accessible to anyone relying on the validity of the certificate authority's certificates.

Public key infrastructure (PKI) is the framework that establishes trust in using public key cryptography to sign and encrypt messages via digital signatures.

3.3 Hashing

As you study this section, answer the following questions:

- What is the output of hashing called?
- What are the five characteristics of a hash function?
- What are some common uses for hashing?
- What type of attack takes advantage of hash collisions?
- What are the main hashing algorithms used?

In this section, you will learn to:

- Use hashes.
- Compare MD5 hashes.

The key terms for this section include:

Term	Definition
Hashing algorithm	A function that converts an arbitrary-length string input to a fixed-length string output. A cryptographic hash function does this in a way that reduces the chance of collisions, where two different inputs produce the same output.
Cryptographic primitive	A single hash function, symmetric cipher, or asymmetric cipher.
Digital signature	A message digest encrypted using the sender's private key that is appended to a message to authenticate the sender and prove message integrity.

Salt	A security countermeasure that mitigates the impact of precomputed hash table attacks by adding a random value to ("salting") each plaintext input.
Key stretching	A technique that strengthens potentially weak input for cryptographic key generation, such as passwords or passphrases created by people, against brute force attacks.
Secure Hash Algorithm (SHA)	A cryptographic hashing algorithm created to address possible weaknesses in multi-domain authentication (MDA). The current version is SHA-2.
Message-Digest Algorithm 5 (MD5)	A cryptographic hash function producing a 128-bit output.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions. • Encryption ○ Algorithms • Hashing • Salting • Digital signatures
TestOut Security Pro	4.0 Data Security 4.2 Implement Encryption Technologies 4.2.1 Encrypt data communications 4.2.2 Encrypt files

3.3.1 Hashing (Lesson Video)

Transcript:

When you're transferring or downloading files, it's important to know that the file is legitimate and hasn't been altered in any way. This is what hashes are for. In this lesson, we'll discuss how hashing works, some common ways they're used, what hashing collisions are, and how to reduce collisions.

Hashing is the process of generating a fixed-length hexadecimal string value from any file type or data.
Let's look at an example. We have an image file that we can run through a hashing algorithm and get a hexadecimal output. This output is the hash. All hashing functions are one-way functions. This means that the hash cannot be reversed. In other words, you can't decipher a hash and find out what the original data was.
A proper hash function has five characteristics. First, it should be deterministic, meaning the same data will always generate the same hash. Next, the hash generation needs to be quick and efficient, and it can't be reverse-engineered.

The hash function should also be collision-resistant and exhibit the avalanche effect, meaning that that if even one tiny bit in the data is changed, the new hash will be completely different.
For example, using the MD5 algorithm, the hash for the message 'Hello world" would look like this. If we capitalize the W and generate a new hash using the same algorithm, we get a very different hash.
There are a few reasons why you may want to generate a hash for some of your data,
such as verifying file integrity, providing digital signatures, and verifying passwords. Let's look at each in more detail.
We download files from the internet all the time. When you do this, you're not always sure that a file is safe and hasn't been tampered with. By providing a hash of the original file, you can verify that the file you downloaded is legitimate.
For example, when an application developer finalizes their program, they can generate a hash of the app and the upload the app to their website along with the generated hash.
Later, this company's website gets hacked, and the hacker replaces the program file with another malicious program by the same name.
When a customer downloads the file, they can generate a hash and compare it to the original hash. Because the app file has been replaced, the hashes won't match, and they'll know that the file has been changed and they shouldn't run it.
Hashing is also used when creating a digital signature for an important message. Using a hash of the data along with our private encryption key, a digital signature is created. This is the equivalent of signing our name to a physical document. When the recipient receives the message, they can verify the digital signature to ensure that the data or message is legitimate.
Hashing is also used to protect passwords. Instead of sending the password in clear text, only a hash of the password is sent. For example, when you use a Microsoft account to sign in to a Windows system, your password is first encrypted using a special hashing algorithm called NTLM. That hash is then checked by Windows, and if it matches, you're signed in.
One of the problems with using hashed passwords is that several online sites have collected massive databases containing a hash for tens of million of possible passwords. Once a hash has been captured, it can be compared with the hashes found in the database, quickly resulting in the password used to create the hash. To keep this from happening, most hashes now also incorporate what's known as a salt.
Salting the hash means that a random number of characters are added to the password before the hash is created.
For example, if the password to be hashed was this, a salt such as this may be added. The string to be hashed becomes this. Since the salt is randomly generated each time, even if the same password is used, and can be varying lengths, it's virtually impossible to create a database containing all the possible salted passwords.
Using hashing helps meet the goals of Information Security by providing file integrity, non-repudiation, confidentiality, and authentication.
Hashing is a very good file verification method, but it's not 100% foolproof. Depending on the algorithm used, there's a potential for hash collisions. A hash collision occurs when two completely different files generate the same hash. A rainbow table attack takes advantage of this weakness. Let's go through the process.
A rainbow table is a table of passwords and their generated hashes. A hacker can use this table to match a captured hash with one in the table. When a match is found, the hacker knows the password, or in the case of a collision, they'll at least know a password that will work.
For example, let's say a user's sign-in password is TestOut, which generates this hash. In the hacker's rainbow table, the password of SecurityPro has also generated the same hash value. This is a hash collision. Taking advantage of this collision, the hacker can gain access to the system using the password of SecurityPro, since it will send the same hash value that the real password uses.
You can reduce the risk of a hash collision by using an algorithm that will generate a longer hash and by salting the hash.
That's it for this lesson. We've discussed how hashing works, including the characteristics of a proper hash function. We reviewed some of the ways hashing is used, including verifying file integrity, digital signatures, and passwords. We ended by discussing hash collisions and how hackers can exploit them using rainbow tables.

3.3.2 Hashing Algorithms (Lesson Video)

Transcript:

Hashing is the process of generating a fixed-length hexadecimal string value from any file type or data. That output is often referred to as the hash or message digest. Hashing can be used to verify file integrity, create digital signatures, provide password verification, and more.

There are different hashing algorithms or functions for different uses. The most common algorithms are MD5 and SHA, and two alternative hash functions are HMAC and RIPEMD. In this lesson, we'll look at these algorithms and discuss how they work.

Message-Digest Algorithm version 5, commonly called MD5, was developed by Ron Rivest in 1991. This hash function generates a 128-bit message digest, but it can't be used for security purposes any longer. One of the key features of a proper hash function is that it's resistant to hash collisions, but the MD5 algorithm is extremely susceptible to collisions. Today, MD5 is mainly used to verify file integrity.

The Secure Hash Algorithm, or SHA, is a family of hashes that are government-standard algorithms and were published by NIST starting in 1993. SHA-2 was published in 2001 and has become one of the standard hash functions in use today. We see it implemented frequently in security communications protocols such as TLS and SSL, PGP, SSH, and IPSec. This algorithm can generate message digests that are 224, 256, 384, or 512 bits in size.

SHA-3 was released in 2015, but it's not meant to replace SHA-2, as it's still a secure hashing method. Because SHA-2 is so heavily integrated in our standard online transactions and communications, it doesn't make sense to switch over yet. SHA-3 uses a completely different process than SHA-2, but the message digests are compatible with SHA-2. What this means for us is that when SHA-2 is inevitably compromised, SHA-3 will be available to take over immediately.

Hash-Based Message Authentication Code, or HMAC, is a type of message authentication code. Like a digital signature, HMAC allows a user to verify that a file or message is legitimate.

When using HMAC, the message sender provides a secret key that's used with a hash function, such as MD5 or SHA, to create a message authentication code. The recipient then uses the key to verify both the integrity and authenticity of the message.

In 1992, a new hash function called RIPE Message Digest, or RIPEMD, was developed based on the MD4 algorithm. In 1996, due to security issues discovered in RIPEMD, a group of Belgian researchers published four updated algorithms. While it's not as popular as SHA-2, the 160-bit function is used frequently with Bitcoin and other cryptocurrencies.

That's it for this lesson. In this lesson, we discussed MD5 and SHA, the most popular hashing functions. MD5 is mostly used for file verification, while SHA-2 is widely used in many security communications protocols. We also discussed two alternative hashing functions, HMAC and RIPEMD. HMAC is used for message or file verification, and RIPEMD is mostly used with Bitcoin and other cryptocurrencies.

3.3.3 Hashing Facts

This lesson covers the following topics:

- Hashing
- Salting and key stretching
- Hashing uses
- Hash collisions
- Hashing algorithms
- Comparing hash values

Hashing

A cryptographic hashing algorithm produces a fixed-length string of bits from an input plaintext that can be of any length. The output can be referred to as a hash or message digest. The function is designed so that it is impossible to recover the plaintext data from the digest (one-way) and that different inputs are unlikely to produce the same output (a collision).

A hashing algorithm is used to prove integrity. For example, Bob and Alice can compare the values used for a password in the following way:

- Bob has a digest calculated from Alice's plaintext password. Bob cannot recover the plaintext password value from the hash.
- When Alice needs to authenticate to Bob, they type their password, convert it to a hash, and send the digest to Bob.
- Bob compares Alice's digest to the hash value on file. If they match, Bob can be sure that Alice typed the same password.

As well as comparing password values, a hash of a file can be used to verify the integrity of that file after transfer.

- Alice runs a hash function on the setup.exe file for their product. They publish the digest on their website with a download link for the file.
- Bob downloads the setup.exe file and makes a copy of the digest.
- Bob runs the same hash function on the downloaded setup.exe file and compares it to the reference value published by Alice. If it matches the value published on the website, Bob assumes the file has integrity.
- Consider that Mallory might be able to substitute the download file for a malicious file. Mallory cannot change the reference hash, however.
- This time, Bob computes a hash that does not match, leading him to suspect that the file has been tampered with.

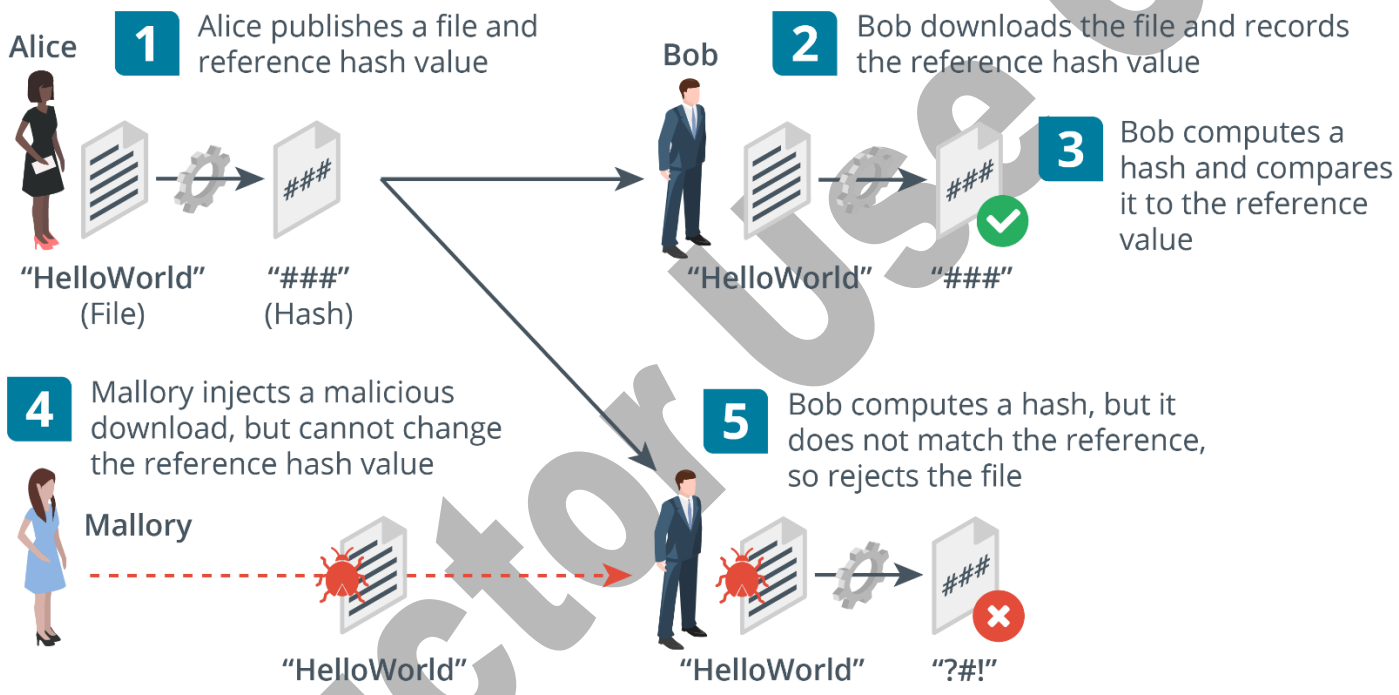

Confirming a file download using cryptographic hashes (Images © 123RF.com)

A single hash function, symmetric cipher, or asymmetric cipher is called a cryptographic primitive . A complete cryptographic system or product is likely to use multiple cryptographic primitives within a cipher suite. The properties of different symmetric/asymmetric/hash types and specific ciphers for each type impose limitations on their use in different contexts and for different purposes.

Encryption can be used to ensure confidentiality. Cryptographic ciphers can also be used for integrity and authentication. If you can encode a message in a way that no one else can replicate, then the recipient of the message knows with whom they are communicating (that is, the sender is authenticated).

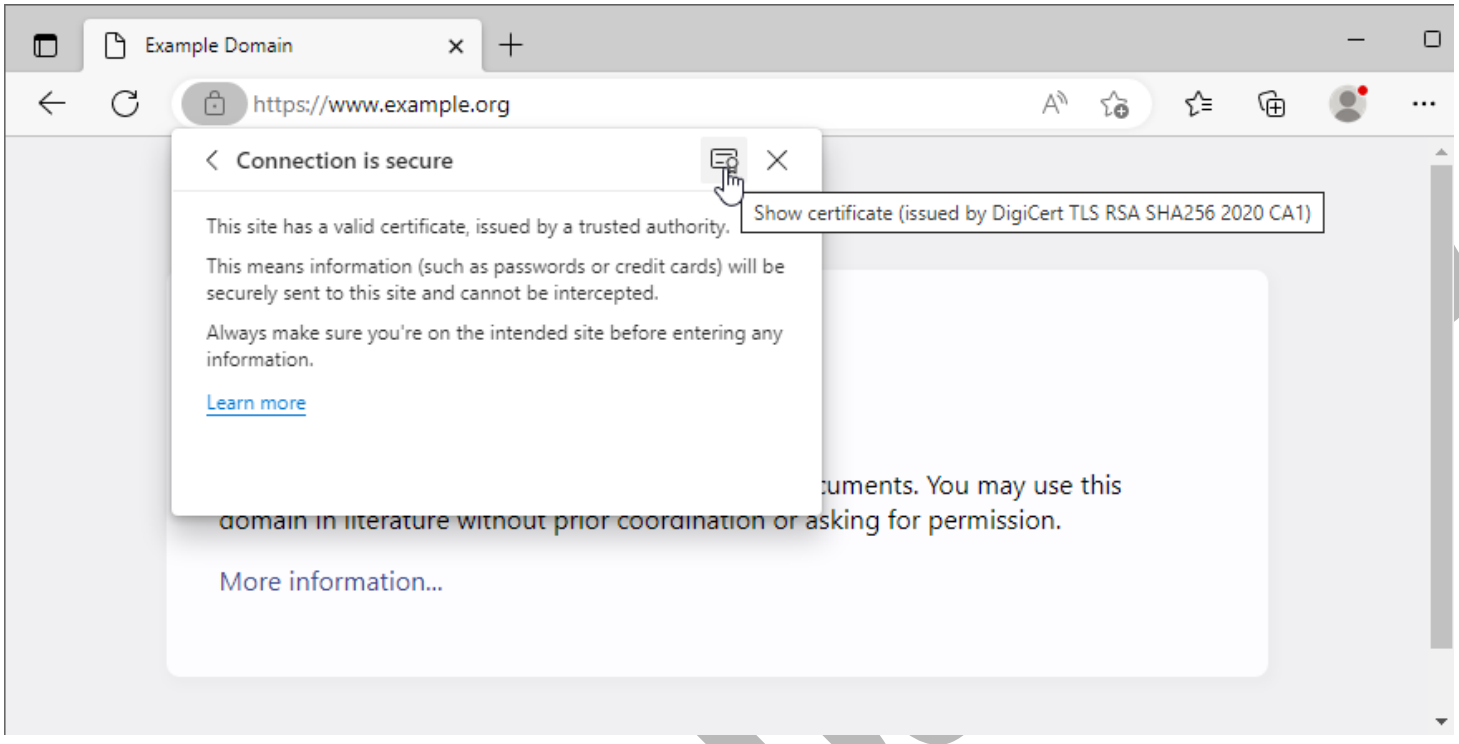

Cryptography allows subjects to identify and authenticate themselves. The subject could be a person or a computer, such as a web server.

Public key cryptography can authenticate a sender because they control a private key that produces messages in a way that no one else can. Hashing proves integrity by computing a unique fixed-size message digest from any variable length input. These two cryptographic ciphers can be combined to make a digital signature :

- The sender (Alice) creates a digest of a message using a pre-agreed hash algorithm, such as SHA256, and then performs a signing operation on the digest using her chosen asymmetric cipher and private key.
- Alice attaches the digital signature to the message and sends both the signature and the message to Bob.
- Bob verifies the signature using Alice's public key, obtaining the original hash.
- Bob then calculates his own digest for the document (using the same algorithm as Alice) and compares it with Alice's hash.

If the two digests are the same, then the data has not been tampered with during transmission, and Alice's identity is guaranteed. If the data had changed or a malicious user (Mallory) had intercepted the message and used a different private key to sign it, the hashes would not match.

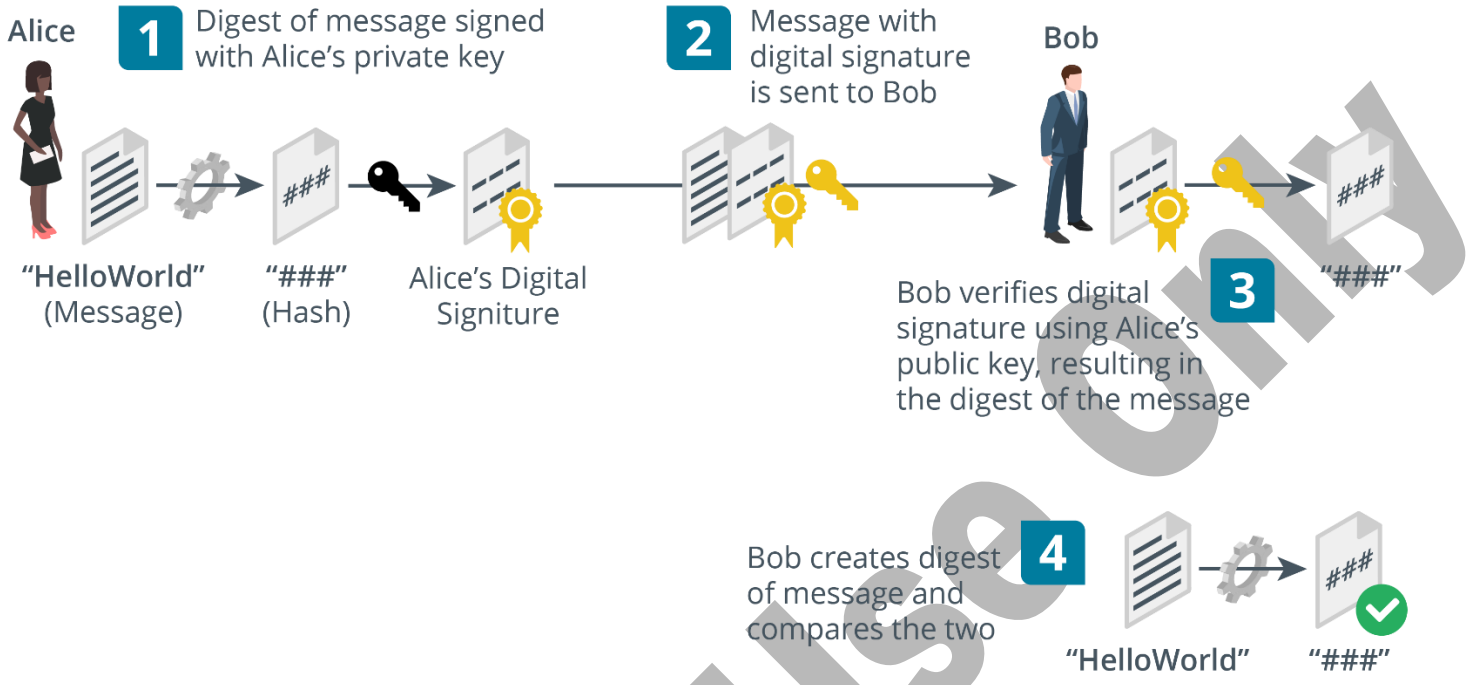

Message authentication and integrity using digital signatures (Images © 123RF.com).

There are several standards for creating digital signatures. The Public Key Cryptography Standard #1 (PKCS#1) defines the use of RSA's algorithm. The Digital Signature Algorithm (DSA) uses a cipher called ElGamal, but the Elliptic Curve DSA (ECDSA) is now more widely used. DSA and ECDSA were developed as part of the US government's Federal Information Processing Standards (FIPS).

Salting and Key Stretching

The values used for a private or secret key must be selected at random. If there is something predictable about the way the value of the key was derived, it has less entropy. Low entropy is a particular concern whenever a cryptographic system makes use of user-generated data, such as a password. Users tend to select low entropy passwords because they are easier to remember. This type of data is too short and ordered to be a good "seed" for key generation. Salting and key stretching help to protect password-derived cryptographic secrets from discovery through cryptanalysis.

Salting

Cryptographic hash functions are often used for password storage and transmission. A hash cannot be decrypted back to the plaintext password that generated it. Hash functions are one way. However, passwords stored as hashes are vulnerable to brute force and dictionary attacks.

A threat actor can generate hashes to find a match for a hash captured from network traffic or a password file. A brute force attack simply runs through every possible combination of letters, numbers, and symbols. A dictionary attack creates hashes of common words and phrases.

Both these attacks can be slowed down by adding a salt value when creating the hash. A salted hash is computed as follows:

(salt + password) * SHA = hash

A unique, random salt value should be generated for each user account. This mitigates the risk that if users choose identical plaintext passwords, there would be identical hash values in the password file. The salt is not kept secret because any system verifying the hash must know the value of the salt. It simply means that an attacker cannot use precomputed tables of hashes. The hash values must be recompiled with the specific salt value for each password.

Key Stretching

Key stretching takes a key generated from a user password plus a random salt value and repeatedly converts it to a longer and more disordered key. The initial key may be put through thousands of rounds of hashing. This might not be difficult for the attacker to replicate, so it does not actually make the key stronger. It does slow the attack down because the attacker has to do extra processing for each possible key value. Key stretching can be performed using a particular software library to hash and save passwords when created. The Password-Based Key Derivation Function 2 (PBKDF2) is widely used for this purpose, notably as part of Wi-Fi Protected Access (WPA).

Hashing Uses

Hashing is often used for the following:

Hash Use	Description
File integrity	Hashes are often used to prove the integrity of downloaded files. When a file is uploaded to a site, a hash can be generated. When the recipient downloads the file, they can create a hash of that file. If the recipient's hash matches the hash of the original file, you know that: • The downloaded file is complete (no missing parts). • The downloaded file was not corrupted during the transfer. • The downloaded file is the same as the original and has not been altered by inserting malicious code or replaced with a virus or malware file. For this reason, files available for download are typically not encrypted, as the hash proves their data integrity.
Digital signatures	A digital signature is a combination of asymmetric encryption and hashing values. A signature provides confidentiality, integrity validation, strong authentication, and non-repudiation. Typically, a digital signature works as follows: 1. A hash value is generated for a message. 2. The hash value is asymmetrically encrypted using the sender's private key. Non-repudiation is provided because only the sender could have encrypted the hash using the private key (only the sender knows the private key). 3. The encrypted hash value and the message are sent. 4. The recipient decrypts the hash using the sender's public key. 5. The recipient hashes the message. 6. Message integrity and sender authenticity (non-repudiation) are confirmed if the two hash values match.
Secure logon credential exchange	Hashes can be used to secure logon credentials during an exchange. The password is used as the key to perform a hash on a text value, and only the hashed value is passed (not the password). The receiving host uses the same method to compare the hashes to verify the identity of the user. Examples of protocols that use this method are: • Challenge-Handshake Authentication Protocol (CHAP)

Hash Use	Description
	• New Technology LAN Manager (NTLM) • Kerberos Passwords can be further secured by salting the hash. This is the process of adding random characters at the beginning or end of the password to generate a completely different hash. If a hacker intercepts the hash, they must also know which portion is the salt before beginning to crack the hash.

Hash Collisions

Hashing is a good file verification method, but it is not perfect. Depending on the algorithm used, there is a potential for hash collisions. A hash collision occurs when two completely different files generate the same hash. Rainbow table attacks take advantage of hash collisions.

- A rainbow table is a table of passwords and their generated hashes. A hacker can use this table to try to match hashes instead of the actual password.
- Hash collisions can be reduced using an algorithm that generates a longer hash and by salting the hash. Salt is random data used as an additional input to the function that hashes data.

Hashing Algorithms

The two most popular implementations of hash algorithms are the Secure Hash Algorithm (SHA) and Message-Digest Algorithm #5 (MD5). Depending on the use, there are different hashing algorithms which can be used. The following table covers some of the more common algorithms.

Hashing Algorithm	Description
Secure Hash Algorithm (SHA)	SHA is a family of hashes. • SHA is a government standard and is considered the strongest algorithm. • First published in 1991 by the National Institute of Standards and Technology (NIST). • SHA-2 was published in 2001 and has become one of the standard hash functions used today. • Used in many security protocols such as TLS, SSL, PGP, SSH, and IPSec. • Generates message digests that are 224, 256, 384, or 512 bits in size. • The longer digests are considered more secure, with the most popular variant being SHA256, which produces a 256-bit digest. SHA-3 was published in 2015 but is not meant to replace SHA-2. SHA-2 has yet to be cracked. NIST wanted an alternative available for people to use. Message digests generated by SHA-3 are fully compatible with SHA-2.

Hashing Algorithm	Description
	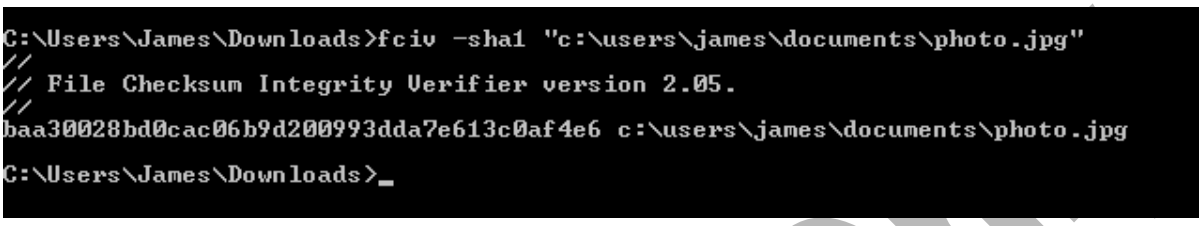 Computing an SHA value from a file (Screenshot used with permission from Microsoft.)
Message-Digest Algorithm 5 (MD5)	MD5 was developed by Ron Rivest in 1991. • MD5 generates a 128-bit message digest. • Many security vulnerabilities have been discovered with MD5. As such, it is no longer viable for security purposes. • MD5 is extremely susceptible to hash collisions. • MD5 is mainly used for file integrity. • MD5 is not considered safe for use, but it might be required for compatibility between security products.
Hash-Based Message Authentication Code (HMAC)	HMAC is a type of message authentication code. Like a digital signature, HMAC allows a user to verify that a file or message is legitimate. • The message sender provides a secret key that is used with a hash function, such as MD5 or SHA, to create a message authentication code. • The recipient then uses the key to verify both the integrity and authenticity of the message.
RACE Integrity Primitives Evaluation Message Digest (RIPEMD)	RIPEMD (RACE Integrity Primitives Evaluation Message Digest, or RIPE Message Digest) is a family of cryptographic hash functions that was first developed in 1992 as part of the EU's RIPE project. The first version was based on the MD4 function. In 1996, in response to security issues discovered in the first version, Belgian researchers developed four updated algorithms. • RIPEMD-128 • RIPEMD-160 • RIPEMD-256 • RIPEMD-320 RIPEMD is not as popular as SHA-2 but is used frequently with Bitcoin and other cryptocurrencies.

Comparing Hash Values

Being able to compare the hash of a file after it has been downloaded to a known good hash helps verify that the file was not altered in transit.

To compare file hashes, you must first generate a hash for the file after it has been downloaded. This can be done using the PowerShell cmdlet named **Get-FileHash** . In the following example, a hash is being generated for the file named Download.zip. The **-a** (or – algorithm) switch specifies the algorithm to use when generating the hash. It is followed by the desired algorithm, in this case, MD5.

Example: **Get-FileHash Download.zip -a md5**

Algorithm	Hash	Path
MD5	A9ESDS4B7288811EC080948E10909A	C: \Downloads\Download.zip

Sample output from Get-FileHash.

Companies often provide a separate file containing the hash of a file that can be downloaded and used to compare to the hash you generate for the downloaded file. This is often in the form of a text file. To view the hash, you can simply open the text file or use PowerShell's **Get-Content** cmdlet to extract the hash from the file.

Example: **Get-Content Download.txt**
Sample Output: **4A84C7958C246E39439C784349F4ZDB4**

Once you have access to the two hashes, you can visually compare them to see if they are the same, or an easier way is to use the **-eq** command to compare to two hash values.

Example: **"39C784349F4ZDB44A84C7958C246E394" -eq "4A84C7958C246E39439C784349F4ZDB4"**

The output will be "True" if the hashes match or "False" if they do not match.

3.3.4 Using Hashes (Demo Video)

Transcript:

In this demonstration, we'll explore how hashing works. I've already downloaded a hashing utility that does both MD5 and SHA-1 hashes. It's called Microsoft File Checksum Integrity Verifier. I've placed it at the root of C: in a folder named FCIV.

As you can see, I have the fciv.exe file, the README.txt that tells you about the file and how to use it, and a little TestOut.txt file that I've created. The name of the executable is obviously fciv.exe, and it uses an MD5 hashing algorithm by default, but it can also use SHA-1.
Let's practice creating a hash of a file. I have this Testout.txt file, and I've entered some text into it. It just says "TestOut." Let's go ahead and close the file.
The first thing we need to do is go to the command prompt. This hashing tool runs off the command prompt. We need to do is get to the correct folder, FCIV at the root of C. Now we're in the right folder, and we can check the directory. You can see the files in the directory, the executable, and the ReadMe. If we open the ReadMe file, it'll show you some of the information about using this file. It shows you some features, and it shows you the syntax for using these features.
Let's go back. We're not going to do anything complicated today. We're going to keep it simple and just create some hashes.
First, we're going to do an MD5 hash. By default, this tool will hash an MD5, and you don't have to state MD5. (You can, but you don't have to.) Let's go ahead and hash our text file. Let's type in â€˜fciv.exe'. Then all we have to do is type â€˜testout.txt'. You can see it's just created our hash for us. The resulting hash is displayed, and MD5 is a 128-bit hashing algorithm, so the output is composed of 128 zeroes and ones. To make it easier to read, the output is then converted from binary to hexadecimal notation. Each hexadecimal character is four bits long, so the hash is composed of 32 hexadecimal characters.

Let's try to modify our Testout.txt file. Suppose that someone has intercepted this file in transit and modified it. Let's just add a period here, at the end. We'll save it and close the file. This isn't much of a change.
Let's say maybe the attacker tried to modify the file without being too obvious. We'll run the hash again. Just press the up arrow and press Enter. We can see that we get a completely different hash value. It's not even close to the original one--it's completely different. If you're using hashing as part of your cryptographic solution, you'll know if data has been modified. The hashes will be marked completely differently.
Let's go back to the file and change it back to the way it was. If we run the hash again, notice the hash is now back to its original value. The fact that it was modified in the past doesn't affect the hash value. The hashing algorithm only evaluates the current content of the data. In this case, the current content is identical to what we had originally hashed, so the hash value is the same. Identical data going into the same hashing algorithm will produce the exact same output.
Let's go ahead and rename the Testout file from Testout to Testout1. We're going to just add a 1 here, and that's it. We have to change this a little bit, just add a 1. Notice that the hash is identical again. The hashing algorithm doesn't care what label you put on the data. It only evaluates the data.
Let's look at an example of SHA, or secure hashing algorithm. Specifically, we're going to use SHA-1. This is similar to MD5. In this tool, we'll use a similar syntax, but we actually have to specify SHA-1.
Let's go back, and we're going to change our file back to TestOut. Instead of entering the fciv.exe, we actually have to specify our SHA-1, because, by default, it does MD5. We have to specify that we want to use SHA-1. Let's go ahead and type in â€˜fciv.exe sha-1 testout.txt'.
All right, so now we can see our file, our SHA-1 hash. Similarly, if we go in and edit our Testout text and add the period, just like we did with MD5, and we save and close, run the same hash again, you can see, we got a different hash, just like with MD5. If we change it back, get rid of the period at the end, save, and run the command again, the same thing occurs as with MD5: we'll get the original hash back.
That's it for this demonstration. This demo provided an overview of how hashing works. We reviewed hashing with MD5 and SHA-1.

3.3.5 Compare an MD5 Hash (Simulation)

Scenario

You are the IT administrator at a small corporate office. You just downloaded a new release for a program you use. You need to ensure the file was not altered before receiving it. Another file containing the original file hash was also downloaded. Both files are located in the C:\Downloads folder.

In this lab, your task is to use MD5 hash files to confirm that the Release.zip file was unaltered.

- Generate a file hash for the Release.zip file.
- View the hash of the original file stored in the **release821hash.txt** file.
- Using the applicable command, compare the original hash of the Release.zip file to its calculated hash to see if they match.

> You can highlight text in PowerShell and right-click it to copy the text to the active line. If using Chromebooks, highlight the desired hash amount and then click on the touchpad using two fingers to copy and paste the value.

- Answer the question.

Explanation

Complete this lab as follows:

1. View the files in the C:\Downloads folder.
    a. Right-click **Start** and select **Windows PowerShell (Admin)** .

b. At the prompt, type **cd C:\downloads** and press **Enter** to navigate to the directory that contains the files.
c. Type **dir** and press **Enter** to view the available files.

2. Confirm that the Release.zip file is unaltered.
   a. Type **get-filehash Release.zip -a md5** and press **Enter** to view the MD5 hash.
   b. Type **get-content release821hash.txt** and press **Enter** to view the known hash contained in the .txt file.
   c. Type ***"new hash"* -eq *"known hash"*** and press **Enter** to determine whether the file hashes match.

   - The new hash is the hash generated by the **get-filehash *file_name* -a md5** command.
   - The known hash is the hash generated by the **get-content *file_name.txt*** command.
   - Include the quotation marks and the file extensions with the file names in the commands.

3. Answer the question.
   a. From the top right, select **Answer Questions** .
   b. Answer the question.
   c. Select **Score Lab** .

3.3.6 Practice Questions (Section Quiz)

q_cryp_hash_collision_secp8

When two different messages produce the same hash value, what has occurred?

Answers:

- Birthday attack
- ***Collision**
- Hash value
- High amplification

Explanation:

A collision occurs when two different messages produce the same hash value.

A birthday attack is a brute force attack in which the attacker hashes messages until one with the same hash is found.

A hash value is the result of a compressed and transformed message (or some type of data) into a fixed-length value.

High amplification means a small change in the message results in a big change in the hashed value.

q_cryp_hash_digest_secp8

Hashing algorithms are used to perform which of the following activities?

Answers:

- Encrypting bulk data for communications exchange.
- ***Creating a message digest.**
- Providing a means for exchanging small amounts of data securely over a public network.

- Providing for non-repudiation.

Explanation:

Hashing algorithms are used to create a message digest to ensure that data integrity is maintained. A sender creates a message digest by performing the hash function on the data files that are transmitted. The receiver performs the same action on the data received and compares the two message digests. If they are the same, the data was not altered.

Symmetric algorithms are used to encrypt bulk data for communications exchange.

Asymmetric algorithms provide a means for exchanging small amounts of data securely over a public network.

Both symmetric and asymmetric algorithms provide non-repudiation.

q_cryp_hash_get-filehash_secp8

You are a system administrator at a software company. You have been tasked with verifying the integrity of a large software update file that has been downloaded from the internet. The software provider has provided a SHA256 hash of the file for verification purposes.

You decide to use the **Get-FileHash** command in PowerShell to calculate the hash of the downloaded file.

Which of the following scenarios best demonstrates the correct use of the **Get-FileHash** command for this purpose?

Answers:

- ***You use the Get-FileHash command to calculate the hash of the downloaded file and compare it with the hash provided by the software provider. If the hashes match, you assume the file has not been tampered with.**
- You use the Get-FileHash command to calculate the hash of the downloaded file and send it to the software provider for verification.
- You use the Get-FileHash command to calculate the hash of the downloaded file and compare it with the hash of the original file on your system. If the hashes match, you assume the file has not been tampered with.
- You use the Get-FileHash command to calculate the hash of the downloaded file and compare it with the hash of another file downloaded from the internet. If the hashes match, you assume the file has not been tampered with.

Explanation:

The proper use of the **Get-FileHash** command is for verifying the integrity of a file. By calculating the hash of the downloaded file and comparing it with the hash provided by the software provider, you can verify that the file has not been tampered with during download.

The software provider has already provided a hash for verification purposes. You should not need to send your calculated hash back to them for verification.

Comparing the hash of the downloaded file with the hash of the original file on your system will not verify the integrity of the downloaded file. The original file on your system may not be the same version as the downloaded file, so their hashes may not match even if the downloaded file has not been tampered with.

Comparing the hash of the downloaded file with the hash of another file downloaded from the internet will not verify the integrity of the downloaded file. The two files are likely to be different, so their hashes should not match.

q_cryp_hash_hash_01_secp8

Which of the following is used to verify that a downloaded file has not been altered?

Answers:

- ***Hash**
- Symmetric encryption
- Asymmetric encryption
- Private key

Explanation:

A hash is a function that takes a variable-length string (message) and compresses and transforms it into a fixed-length value. Hashes ensure the data integrity of files and messages in transit. For example, when users post files for download, they often create a hash value for the file. After you download the file, you can create a hash using the same algorithm. If the hash values match, you know that the file you have matches the original file.

Symmetric encryption is typically used for fast data encryption.

Asymmetric encryption is used for encrypting small amounts of data or exchanging keys used with symmetric encryption.

A private key is one of the keys used in asymmetric encryption.

q_cryp_hash_hash_02_secp8

You have just downloaded a file. You create a hash of the file and compare it to the hash posted on the website. The two hashes match.

What do you know about the file?

Answers:

- ***Your copy is the same as the copy posted on the website.**
- You can prove the source of the file.
- No one has read the file contents as it was downloaded.
- You are the only one able to open the downloaded file.

Explanation:

A hash is a function that takes a variable-length string (message) and compresses and transforms it into a fixed-length value. Hashes ensure the data integrity of files and messages in transit. The sender and the receiver use the same hashing algorithm on the original data. If the hashes match, it is assumed that the data is unmodified.

Hashes do not ensure confidentiality (in other words, hashes are not used to encrypt data).

Non-repudiation proves the source of a file and is accomplished using digital signatures.

q_cryp_hash_hash_04_secp8

A cybersecurity analyst working for a financial institution handles sensitive customer data and must ensure its integrity and confidentiality. Part of the analyst's responsibilities is to evaluate the hashing algorithms used to protect the data's integrity.

Which statement accurately describes the purpose and functionality of hashing algorithms in data security?

Answers:

- ***Hashing algorithms generate a fixed-size output (a hash value), which represents the input data and is unique to each input.**
- Hashing algorithms encrypt data to make it unreadable to unauthorized users.
- Hashing algorithms use a public and private key pair to transmit data over the network securely.
- Hashing algorithms compress data to reduce storage requirements and improve performance.

Explanation:

Hashing algorithms generate a fixed-size output called a hash value, which represents the input data and is unique to each input. Additionally, they generate unique, fixed-size output (hash value) for any given input. The hash value is a digital fingerprint of the input data and is unique to each input, ensuring data integrity and allowing for efficient data comparison.

Hashing algorithms do not relate to encryption but rather to data integrity verification.

Hashing algorithms do not involve the use of public and private key pairs.

Hashing algorithms do not relate to data compression; instead, they create a unique representation of the input data.

q_cryp_hash_hmac_secp8

Which of the following is a message authentication code that allows a user to verify that a file or message is legitimate?

Answers:

- SHA
- MD5
- RIPEMD
- ***HMAC**

Explanation:

Hash-based Message Authentication Code (HMAC) is a type of message authentication code. Like a digital signature, HMAC allows a user to verify that a file or message is legitimate.

SHA is a family of hashes that is used in many different security protocols.

MD5 was developed in 1991 and is no longer viable for security purposes.

RIPEMD is a family of cryptographic hash functions that was first developed in 1992 as part of the EU's RIPE project.

q_cryp_hash_rainbow_secp8

An attacker is attempting to crack a system's password by matching the password hash to a hash in a large table of hashes he or she has.

Which type of attack is the attacker using?

Answers:

- Cracking
- Brute force
- ***Rainbow**
- RIPEMD

Explanation:

A rainbow attack uses rainbow tables. A rainbow table is a table of passwords and their generated hashes. A hacker can use this table to try to match hashes instead of the actual password.

Cracking is the process of finding a password.

A brute force attack does not use a table of hashes.

RIPEMD is a family of cryptographic hash functions that was first developed in 1992 as part of the EU's RIPE project.

q_cryp_hash_ripemd_secp8

What is the primary use of the RACE Integrity Primitives Evaluation Message Digest (RIPEMD)?

Answers:

- It is primarily used for file compression.
- ***It is primarily used in Bitcoin and other cryptocurrencies.**
- It is primarily used for creating digital watermarks.
- It is primarily used for email encryption.

Explanation:

RIPEMD, specifically the RIPEMD-160 variant, is used in the creation of Bitcoin wallet addresses. It is part of the process that converts a user's public key into a Bitcoin address.

RIPEMD is a cryptographic hash function, not a compression algorithm. It is used to ensure data integrity and authenticity, not to reduce the size of data files.

While cryptographic hash functions can be used in the creation of digital watermarks, this is not the primary use of RIPEMD. Digital watermarking typically involves embedding information into a digital signal, which is not the function of RIPEMD.

While cryptographic hash functions are used in various aspects of email encryption, RIPEMD is not specifically designed or primarily used for this purpose. Other protocols and algorithms, such as RSA and AES, are more commonly used for email encryption.

q_cryp_hash_salting_secp8

What is the process of adding random characters at the beginning or end of a password to generate a completely different hash called?

Answers:

- ***Salting**
- Deterministic
- Collision

- Avalanche

Explanation:

Salting is the process of adding random characters at the beginning or end of the password to generate a completely different hash. If a hacker intercepts the hash, he or she would need to also know which portion is the salt before beginning to crack the hash.

Deterministic is a characteristic of a hash function that means the same data always generates the same hash.

A collision is when two different pieces of data generate the same hash.

The avalanche effect states that changing any bit of data results in a completely different hash.

q_cryp_hash_secure_logon_secp8

You are a security analyst at a large corporation. The corporation is implementing a new system that requires secure logon credential exchange between different departments.

The corporation decides to use a cryptographic hashing algorithm for this purpose.

Which of the following scenarios best demonstrates the correct use of hashing for secure logon credential exchange?

Answers:

- ***Each department calculates a hash of their password and sends it to the other departments. The receiving department compares this hash with the hash of the sending department's password they have on file. If the hashes match, the receiving department assumes the sending department has authenticated itself.**
- Each department sends their password to the other departments. The receiving department calculates a hash of the received password and compares it with the hash of the sending department's password they have on file. If the hashes match, the receiving department assumes the sending department has authenticated itself.
- Each department calculates a hash of their password and sends the actual password along with the hash to the other departments. The receiving department then verifies the password by comparing it with the hash.
- Each department shares their passwords with each other, calculates the hash of their own password, and if the hashes match, they assume they have authenticated each other.

Explanation:

The proper use of a hashing algorithm for secure logon credential exchange includes each department calculating a hash of their password and sending it to the other departments. The receiving department compares this hash with the hash of the sending department's password they have on file. If the hashes match, the receiving department can be sure that the sending department has authenticated itself.

Departments should never send their actual passwords, even if the receiving department is using hashing for authentication. The purpose of using hashing is to avoid sending the actual password.

Departments should never send their actual passwords, even if they are also sending the hash. The purpose of using hashing is to avoid sending the actual password.

Departments should never share their actual passwords with each other, even if they are using hashing for authentication.

q_cryp_hash_weak_secp8

Which of the following is no longer valid for security purposes?

Answers:

- ***MD5**
- SHA-1
- DES
- AES

Explanation:

MD5 is the weakest hashing algorithm. It produces a message digest of 128 bits. Many security vulnerabilities have been discovered with MD5. As such, it is no longer viable for security purposes.

SHA-1 is more secure because it produces a 160-bit message digest.

Both DES and AES are symmetric encryption algorithms. DES is weaker than AES.

q_md5_hash

Do the file hashes match?

Answers:

- No
- No
- Yes
-

3.4 Encryption

As you study this section, answer the following questions:

- Which editions of Windows include Encrypting File System (EFS)?
- Why would you create a Data Recovery Agent (DRA)?
- Which standard does Pretty Good Privacy (PGP) and GNU Privacy Guard (GPG) follow?
- What partition/volumes are created when implementing BitLocker?
- What are three methods of database encryption?

In this section, you will learn to:

- Encrypt files using EFS.
- Encrypt files using GPG.
- Configure BitLocker with a Trusted Platform Module (TPM).

The key terms for this section include:

Term	Definition
Data at rest	Information that is primarily stored on specific media, rather than moving from one medium to another.
Data in transit (or data in motion)	Information that is being transmitted between two hosts, such as over a private network or the internet.
Data in use (or data in processing)	Information that is present in the volatile memory of a host, such as system memory or cache.
Transport/communication encryption	Encryption scheme applied to data-in-motion, such as WPA, IPsec, or TLS.
Key exchange	Any method by which cryptographic keys are transferred among users, thus enabling the use of a cryptographic algorithm.
Hash-based Message Authentication Code (HMAC)	A method used to verify both the integrity and authenticity of a message by combining a cryptographic hash of the message with a secret key.
Full disk encryption (FDE)	Encryption of all data on a disk (including system files, temporary files, and the page file) that can be accomplished via a supported OS, third party software, or at the controller level by the disk device itself.
Self-encrypting drives (SED)	A disk drive where the controller can automatically encrypt data that is written to it.
Key Encryption Key (KEK)	In storage encryption, the private key that is used to encrypt the symmetric bulk media encryption key (MEK). This means that a user must authenticate to decrypt the MEK and access the media.
Opal Storage Specification	Standards for implementing device encryption on storage devices.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions.
	• Encryption
	  ○ Full-disk
	  ○ Database

	• Tools ○ Trusted Platform Module (TPM) 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. • Encryption
TestOut Security Pro	4.0 Data Security 4.2 Implement Encryption Technologies 4.2.1 Encrypt data communications 4.2.2 Encrypt files

3.4.1 Encrypting File System (Lesson Video)

Transcript:

Microsoft's Encrypting File System, or EFS, was introduced with NTFS version 3 and has been included in every version of Windows since Windows 2000, except for in the Home editions.

EFS makes it simple for users to encrypt files on their Windows systems.
In this lesson, we'll look at the EFS encryption process, recovering data, and some security considerations with EFS.
EFS combines the speed of symmetric encryption with the convenience of asymmetric encryption.
However, keep in mind that EFS is used to encrypt individual files or folders on the system, not the entire drive. You would have to use BitLocker for that.
When you want to encrypt a file, Windows generates a pseudo-random number called the File Encryption Key, or FEK.
Windows uses the FEK with the AES encryption algorithm to encrypt the file.
To further protect everything, the FEK is then encrypted using the user's public key and stores the encrypted FEK in the file's header in a special location called the Data Decryption Field, or DDF.
The decryption process is just the opposite.
The user's private key is used to first unlock the DDF and get the FEK. The FEK is then used to decrypt the file.
By combining the security of the symmetric keys and the convenience of the asymmetric keys, Microsoft has made it easy for users to encrypt their data.
By default, the only user that can decrypt files is the user who encrypted them. If that user account becomes corrupted somehow or the password is forgotten, any encrypted files are lost. To prevent this, we need to setup a Data Recovery Agent, or DRA.
The Data Recovery Agent, or DRA, is simply another account that can decrypt data that's been encrypted by other users on a Windows operating system.
In older versions of Windows, the system administrator was automatically configured as the DRA, but in newer versions, the DRA is not automatically defined.
Instead, you have to go into Group Policy on the local computer to setup the DRA.
When working on a Domain network, the DRA is defined in Active Directory.
Keeping the decryption keys safe is vital to protecting your data and being able to access it when needed.
To do this, you can backup the decryption keys to a USB drive. Then, if something catastrophic happens to the Windows system, you still have access to the files, you'll be able to decrypt the files using the backup of the decryption keys.
When using EFS, there are some security issues you need to keep in mind, such as what happens when files are moved.
If you move or copy an encrypted file to a different location on the same partition, or to another NTFS partition, your file will stay encrypted.
However, if you move or copy a file to a FAT based partition, the file will be decrypted automatically, as these file systems don't support encryption.

We need to be careful with this because by default, USB drives are formatted with the FAT32 or exFAT file systems as these are most compatible with other operating systems.
When you copy an encrypted file over, you won't be notified that the file is decrypted. This can lead to some serious security issues.
The other security concern is that the user's private key is protected only by their password. If the user has a weak password and the system is compromised, all encrypted files are vulnerable. It's absolutely vital that users have strong passwords and follow proper password security protocols.
That's it for this lesson.
In this lesson we discussed EFS and how the encryption process works.
We then looked at why a Data Recovery Agent should be created so there's a second account that can access encrypted files in case something happens.
And finally, we reviewed some security considerations with EFS including moving encrypted files and using strong passwords.

3.4.2 Encrypt Files (Demo Video)

Transcript:

In this demonstration, we're going to talk about protecting sensitive files with encryption using the Encrypting File System, or EFS, on Windows. We'll demonstrate with files that are stored on an NTFS file system volume. Let's get started.

We begin by going down here to File Explorer. I'm going to go to my C: drive, and you can see that I've created this folder called Confidential. I have two different files in here.
The first one is my Org Chart, and the second one is where I keep my performance reviews. Now, the Org Chart file isn't really sensitive at all. Pretty much everybody in the organization knows who reports to who.
The Performance Reviews file, on the other hand, is highly confidential. I don't want any of my employees to access it. One thing I can do to prevent this is to encrypt the file so that only my user account can open it.
To do this, I right-click the file, go to Properties, and then to Advanced. I can either compress the file's contents or encrypt it. Notice that you can't select both at the same time. We're just going to select Encrypt contents to secure data, click OK, and click Apply.
We're prompted to either encrypt the file and its parent folder, which would encrypt the entire Confidential folder and everything in it, or the file alone. In this case, we want to select Encrypt only the file. Select OK and then Ok again. You can see that, right here, there's a little lock on the icon now. That means that the file is now fully encrypted.
To clarify, I'm logged in as the TestOut user right now. At this point, only this account is allowed to access this file's contents. If I double-click it, I can easily open it up and see everything. I read that all my employees have met expectations. That means I have some really great employees.
Let's move on to testing what we've done by logging on as a different user—Kim Sanders—whose user account is named ksanders.
Now, ksanders is a member of my administrative group, and so she's obviously an administrator on this system. Let's go ahead and go back to the file system the same way we did before—File Explorer > C: drive > Confidential.
We see the encrypted Performance Reviews file here. Let's try to open it. You can see that access to this file has been denied to this user because, although she's a system administrator, she doesn't have access to this specific file.
This file wasn't set up with NTFS permissions that are blocking the access. It's the encryption. Since ksanders is an administrator, she technically has access to all the data that other users have on the system. Let's go ahead and close this window.
By default, only the user who encrypted the file is allowed to decrypt it and access its contents, but we can change that.
We're going to go ahead and sign out of ksanders and log back in to our TestOut user account again.
We'll navigate again to File Explorer > C: drive > Confidential > Performance Reviews. Right-click, go to Properties, and click the Advanced button. Here you can see that the Details option is now available, whereas before we encrypted the file, it wasn't.
If we go into Details, we can specify other users that we want to give file access to. Let's click Add. Notice that ksanders isn't listed here. It just shows my TestOut user.
The issue is that my TestOut user has a certificate. It was created automatically when I encrypted this file. The ksanders user, however, doesn't have a certificate.

That means that before we can add the ksanders user account to the list of users who are allowed to access the Performance Reviews file, we have to create a certificate for her, which we do in the ksanders profile. So let's go back out again and switch our user. Sign out. Now log back in as ksanders and provide the password.
With the Start button, search for Control Panel > User Accounts > User Accounts again. Here, we'll go to Manage your file encryption certificates. This wizard easily walks us through creating a file encryption certificate. We just need to hit Next. Then we click Create a new certificate and Next again.
We have some options here. The basic one is to create a self-signed certificate. If you had a smart card or if your domain had a certificate authority, you could use either of those as well. Click Next. We're prompted to back up the certificate after it's created. That's a really good idea. Let's say that something were to happen and we needed to re-create the ksanders user account. The new key won't be the same, and ksanders will lose access to all the files that she encrypted on the system.
But if you have a backup of the original encryption key, that'll be the one that's used to encrypt the file. Instead of creating a new key, you just restore your old one from backup. You won't lose access to anything.
We won't worry about that specifically today, but just know that this is important to think about. I'll just click Back up later and Next to move on. We need to check the I'll update my encrypted files later box and click Next as well before we can keep going.
You now see that our certificate is created. Click Close. Now what we need to do is go back in to our TestOut user account and allow Kim Sanders access via that certificate.
Let's close this and go back to Sign out. Now we're going to log back in as our TestOut user once again. Let's go back to our file system and to the Confidential folder.
We'll return to the Performance Reviews file, click Properties, go to Advanced, and then to Details. We need to click Add, select More choices, and then you can see the ksanders certificate here. Just select it and click OK.
We're warned here that the certificate is self-signed, and it's not issued by a trusted root certificate authority. It wants to know if we trust it. Since we created it, we should obviously trust it. I click OK, OK again, and Apply.
And now, one last time, we're going to sign out and back in to the ksanders user account. Select Kim Sanders and back through File Explorer > C: drive > Confidential > Performance Reviews and double-click. As you can see, the file opened without a problem.
Now ksanders has access to an encrypted file that was encrypted by another user.
So that's it for this demonstration. In this demo, we talked about EFS file encryption on Windows. We first talked about how this works with the NTFS file system. Then we encrypted a sensitive file for the TestOut user account. After that, we switched users and created a certificate for the ksanders user account, and we assigned that account access to the encrypted file with the security certificate.

3.4.3 Encrypt Files with EFS (Simulation)

Scenario

You share a computer with other users at work. You want to secure the contents of the Finances folder so that unauthorized users cannot view its contents.

In this lab, your task is to:

- Encrypt the **D:\Finances** folder and all of its contents.
- Give John file access to the encrypted **D:\Finances\2023report.xls** file by adding the encryption certificate.

Explanation

Complete this lab as follows:

1. Open the D:\ drive.
    a. From the Windows taskbar, select **File Explorer** .
    b. From the left pane, select **This PC** .
    c. From the right pane, double-click **Data (D:)** .
2. Encrypt the Finances folder.

a. Right-click **Finances** and then select **Properties** .
b. Select **Advanced** .
c. Select **Encrypt contents to secure data** and then select **OK** .
d. Select **OK** to close the properties dialog.
e. Select **OK** to confirm the attribute changes.

3. Give John authorization to modify the encrypted 2023report.xls file.
   a. Double-click **Finances** .
   b. Right-click **2023report.xls** and then select **Properties** .
   c. Select **Advanced** .
   d. Select **Details** .
   e. Select **Add** .
   f. Select **John** and then select **OK** .
   g. Select **OK** as many times as needed to close all remaining dialogs.

3.4.4 PGP and GPG (Lesson Video)

Transcript:

When you're encrypting files or emails, you need the help of a utility. Windows automates this process with either BitLocker or EFS. In a Unix-based operating system, such as Linux or Apple's OS X, we can use GNU Privacy Guard, or GPG. GPG is based on an older utility, Pretty Good Privacy, or PGP. In this lesson, we'll cover how both utilities work and how to use them.

PGP is an encryption program first developed in 1991 based on the OpenPGP standard. PGP combines the use of symmetric and asymmetric keys and can be used to send encrypted messages and encrypt data.
To encrypt data, PGP generates a large, random one-time use session key that's used for encryption. The session key is then encrypted using the receiver's public key, and both are combined to send the encrypted message.
When the receiver gets the message, they use their private key to decrypt the session key, which is then used to decrypt the message.
Even though PGP is an old utility, it's still considered the standard for encrypting messages, because at the time this video was recorded, it's never been cracked. PGP was purchased a while ago and commercialized. It's owned by NortonLifeLock, formally known as Symantec, and provides products that can protect all sorts of devices, even smartphones.
In response to PGP becoming a commercial product, GNU Privacy Guard was created in 1999. GPG is a command line utility that's used to encrypt and decrypt data and messages. GPG functions just like PGP. It uses both symmetric and asymmetric keys to encrypt and secure data and messages.
To generate a random session key, the user performs actions on the computer, such as typing on the keyboard or moving the mouse. This helps to ensure that the key is truly random.
Because it's an open-sourced utility, GPG can be used on many different systems, including Windows, Linux, Android, and Apple's OS X.
That wraps up this lesson. In this video, we looked at two popular utilities for encrypting files and messages, Pretty Good Privacy and GNU Privacy Guard. Both utilities use a hybrid cryptographic model and are very secure. The biggest difference between them is that PGP is a commercial product, and GPG is a free open-source utility.

3.4.5 Encrypting Files with GPG (Demo Video)

Transcript:

In this demonstration, we're going to review how to encrypt files on a Linux system using GPG. The GPG utility is the open-source equivalent of the PGP utility.

First, we need to verify that GPG's installed on the system, so we're going to type â€˜gpg â€“version'. As you can see here, it'll show us the version information. It wouldn't show this if GPG wasn't installed. Most Linux distributions should have this installed by default. It's currently installed on the system. Let's go ahead and clear our screen.

Now we need to generate a key pair. To do that, we're going to type â€˜gpg --gen-key'. If you want to have more customization, you can use the full-featured key generation dialog as listed. It wants to know the real name. We're just going to put our username in this field, rmckay. Our email address is rmckay@testout.com. This last menu wants us to verify the information. We'll okay it.
Another dialog box will come up asking us to enter a passphrase. For this demo, I'm just going to use the word â€˜TestOut'. This gives us a warning, saying it's an unsecure passphrase. Normally, you would make this password more secure, but for now, we're going to leave it as-is. We'll re-enter the same passphrase. Now our public and secret keys are created.
We're going to encrypt an example file in our Home directory. It's called example.txt. We can see this by doing an â€˜ll' to list the directory.
Now we want to go ahead and encrypt our file using GPG. You don't have to specify the entire directory path of the file if you're in the directory where the file exists.
To encrypt, we're going to type in â€˜gpg -e example.txt' and then the name that we created, â€˜rmckay'. Then we're going to click Enter twice, as there's no more information we need to add.
Now, if we type â€˜ll' and look, we can see the original example.txt, and we can see the encrypted version of example.txt.gpg. If we use the CAT command to view this encrypted version of the file, you can see that it's all scrambled. Clearly, we can't read the file or tell what it says. At this point, obviously, we need to decrypt this file before we can read it. That's what we're going to do next. I'm just going to clear the screen again.
Now, we're going to type â€˜gpg â€“output'. We're going to follow this with the file name for the new decrypted file that we're going to create, â€˜example.txt.decrypted'. Okay. Now we're going to do â€˜--decrypt example.txt.gpg'.
In this case, we're decrypting with the same user we created a public and private key with, so it won't prompt us for a password. Okay. Now we can use â€˜ll' to list the files. We see our decrypted file right here. If we use the CAT command, we see what's inside example.txt.decrypted. This verifies that the GPG utility successfully decrypted the file.
That's it for this demonstration. We discussed how to encrypt and decrypt files in a Linux system using GPG.

3.4.6 BitLocker and Database Encryption (Lesson Video)

Transcript:

Organizations often store sensitive data on devices and in databases. Implementing proper data encryption is key to securing this sensitive data. Using BitLocker and proper database encryption can help protect data if a physical device, such as a laptop, is stolen or if a hacker gains access to a database.

A lost or stolen computer can be catastrophic to an organization if it holds confidential information. With the release of Windows Vista, Microsoft introduced BitLocker to address this concern. BitLocker is a powerful encryption tool that, instead of encrypting individual files and folders, encrypts an entire volume, including operating system files. BitLocker is designed to protect data from unauthorized access, even if the drive is moved to another computer.
It's important to note that BitLocker isn't available on Home editions of Windows.
To implement BitLocker on a computer, the hard disk must be partitioned with two volumes. The system, or boot, volume contains the Windows boot files and is created during Windows installation. The standard volume contains all other data. The system volume won't be encrypted, but the standard volume will be.
One of the newer features of BitLocker in Windows 10 is the ability to only encrypt used space. This makes the encryption process so much faster. Previously, the entire drive, even space not in use, was encrypted. This process could take hours depending on the size of the drive.
BitLocker can also use the computer's Trusted Platform Module, or TPM, chip to verify the integrity of the system's boot files as long as the chip is at least version 1.2. It does this by encrypting the boot files and stores the encryption key in the TPM chip. When you log in to Windows, BitLocker automatically unlocks the encrypted drive. If the drive is moved to another computer, the encryption key won't match up, and the drive can't be accessed. It's possible to configure BitLocker to protect the system files without having a TPM chip, but you have to insert a startup USB key or have a system volume password enabled to boot into Windows.
When BitLocker is enabled, Windows creates the recovery key. This randomly generated key will be used if the hard drive needs to be moved to a different system, if changes are made to the startup files, or if BitLocker goes into a locked state and needs to be accessed. The recovery key is different from the user-generated password created during the configuration process.

Obviously, this is an incredibly important key. Windows gives you the option to save the key to your Microsoft Account, a USB flash drive, or a file on the local computer. You can also print the key out, and if you're on a domain, you can store the key in Active Directory. It's a good idea to back up this key multiple ways and then store it safely locked.
If you're on a domain, there are additional options for recovering data if the user password is lost. Using the stored recovery key would be the first and easiest option, especially if the key was stored in Active Directory. If the key can't be found, a Data Recovery Agent, or DRA, can be used. The DRA is just another account that has rights to decrypt the drive. This option must first be configured in Group Policy. If the hard drive contains the OS files, it'll need to be installed in another system as a data drive before the DRA can decrypt it.
When implemented properly, BitLocker is a powerful tool that can be used to protect data in case a laptop or other device that contains sensitive information is stolen or lost.
Many organizations store important sensitive data, such as customer billing information, in databases. Keeping this data encrypted helps protect it if a hacker ever gains access to the database.
The three main methods of database encryption are transparent, column-level, and application-level. Before we look at these methods in detail, let's review the structure of a database.
Databases are made up of multiple tables that use columns and rows filled with data. For example, we might have a table labeled Customers. In that table, we have columns labeled Name, Number, Email, and Address. The rows in each column are filled with the pertinent information.
Transparent data encryption, or TDE, encrypts the entire database and all backups. TDE is used for data at rest--data that's not in current use. This method is called transparent encryption because when an authorized user needs to access the data, it's automatically decrypted, so the user doesn't see the process or need to do anything extra.
Column-level encryption allows the administrator to encrypt each column using different keys. This increases security because multiple encryption keys are required to access all of the data. Keep in mind that this method does slow the database's performance.
In application-level encryption, the program that's used to create or modify the data is responsible for encrypting the data. This works well because the data is encrypted before it even hits the database. The drawback of this method is that the amount of resources required for setup can be prohibitive.
When an organization maintains sensitive data, they need to ensure that it's kept safe. Any of the methods in this video will help keep databases secure.
That wraps up this lesson. We've covered some solutions for encrypting sensitive data. BitLocker can be used to encrypt an entire hard drive or volume, which will prevent access if a device is stolen. And there are several methods an organization can use to encrypt sensitive data that's stored in databases.

3.4.7 Configuring BitLocker (Demo Video)

Transcript:

In this demonstration, we're going to show you how to use BitLocker to encrypt a drive on Windows. Before anything else, you should know that BitLocker isn't available unless you have Windows Professional version or higher. Home version doesn't support this feature. That being said, let's go to the Start menu and type bitlocker. Clicking Manage BitLocker brings up the settings we want. Now, the first thing we have to cover in order to understand the BitLocker setup is a TPM.

The TPM, or Trusted Platform Module, is a chip on the motherboard that's known as a cryptographic module. Basically, it's what stores your BitLocker keys. Without the TPM, you'd have to enter a long recovery key each time, especially if the BitLocker drive were installed on a computer other than the original one. It used to be common for motherboards to ship without a TPM, but as security has progressed, they're on pretty much every motherboard nowadays. In fact, a TPM is actually required with the latest versions of Windows.
Okay, let's move on. Look at the TPM Administration down here. This is an area that I'd highly suggest checking first before you consider BitLocker for your system. The TPM needs to be in a ready for use status and enabled in the UEFI or BIOS. If it was previously used for BitLocker, you may have to clear the TPM chip before using it. This system we have here hasn't been encrypted before, so we don't need to do that. Also, if it isn't in a ready for use status, you'll need to use this Prepare the TPM option up here. Let's close this Window.
To start the process, we need to click Turn on BitLocker. First, it'll ask us how we want to back up our recovery key. You have a few options, but we're just going to use Save to a file. When saving to a file, you're supposed to put this key on something else other than the drive being encrypted, usually something like a flash drive or network drive. We have a

flash drive as drive letter E: that we can use to store this. After selecting the drive and clicking Save, it tells us that our recover key has been saved. Were able to select Next.
Here we have options to encrypt just the used disk space or encrypt the entire drive. Typically, the Encrypt whole drive option takes longer because it doesn't care if space is used or not. This PC happens to be new, so we'll just choose Encrypt used disk space only. Okay, click Next.
Now, Microsoft introduced a new type of encryption method with the most recent versions of Windows. We could use it, but for now, we'll just leave it set to New encryption mode. Click Next.
The last part asks if you're ready to encrypt. I like to select the Run BitLocker system check just in case something ends up causing problems. In most cases, BitLocker deployment should go fine, but there are instances where it doesn't work properly. It's a good idea to have your computer backed up in some way so that your data can't be lost forever during this process. Click Continue.
Now we're ready to restart the computer. You'll get a pop up that lets you know a restart is required before you can keep going with the encryption. I'm going to reboot this computer and pause the recording.
Alright, after the restart, we see an option down in our taskbar that shows us that the drive is encrypting. When we click on it, it gives us a percentage of where it's at. Let's pause the recording again while this finishes.
Great! Our encryption is now done, and we're ready to go. Just to show you some options you have after the encryption is complete, let's go back. You can click Manage BitLocker, or you can access it like before through the Start menu.
The first option we have is Suspend protection, which means it'd put BitLocker on hold from securing your system. This might be good if Windows updates are needed and your PC is already secured with a BitLocker PIN. Make note that the BitLocker PIN won't be enabled unless you modify Group Policy and turn it on. That adds an extra layer of security, so it's definitely something to consider. Now, this isn't the same as a recovery key, as the recover key is typically 48 digits and is used as a last resort to decrypt your drive. Backup your recover key gives us the same box we had before, remember the Azure AD, Save to file, and Print options from earlier? And the last one is Turn off BitLocker. This process will take nearly the same amount of time to decrypt since it has to undo all the changes made during encryption.
That's it for this demonstration on BitLocker. In this video, we talked about how to encrypt a system drive. First, we looked at the TPM and the purpose that it serves in the BitLocker process. Then we discussed which settings you have available after BitLocker has successfully encrypted a drive.

3.4.8 Configure BitLocker with a TPM (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. The employee in Office 1 is working on a very sensitive project. Management is concerned that if the hard drive in the computer were stolen, sensitive information could be compromised. As a result, you have been asked to encrypt the entire System volume. The Office 1 computer has a built-in TPM on the motherboard.

In this lab, your task is to configure BitLocker drive encryption as follows:

- From within the computer's BIOS, turn on and activate TPM Security.
- From Windows, turn on BitLocker for the System (C:) drive.
- Back up the recovery key to the **\\CorpServer\BU-Office1** folder.
- Encrypt the entire System (C:) drive.
- Use the new encryption mode.
- Run a BitLocker system check.

Explanation

Complete this lab as follows:

1. (Optional) Try to enable BitLocker.
    a. From the taskbar, select **Search** .
    b. In the Search field, type **Control**
    c. From *Best match* , select **Control Panel** .

d. Select **System and Security** .
e. From the right pane, select **BitLocker Drive Encryption** .
f. Under *Operating system drive* , select **Turn on BitLocker** .
An error message at the bottom of the screen indicates that a TPM security device was not found.
g. Select **Cancel** .

2. Access the BIOS settings.
a. Right-click **Start** , then select **Shut down or sign out** > **Restart** to reboot your computer.
b. When the TestOut logo appears, press **Delete** or **F2** to enter the BIOS.
3. Turn on and activate the TPM.
a. From the left pane, expand and select **Security** > **TPM Security** .
b. Select **TPM Security** and then select **Apply** .
c. Select **Activate** and then select **Apply** .
d. Select **Exit** .
Your computer will automatically reboot.
4. Turn BitLocker on.
a. From the taskbar, select **Search** .
b. In the Search field, type **Control**
c. From *Best match* , select **Control Panel** .
d. Select **System and Security** .
e. Select **BitLocker Drive Encryption** .
f. Under *Operating system drive* , select **Turn on BitLocker** .
Windows begins the Drive Encryption setup.
5. Back up a BitLocker recovery key.
a. Select **Save to a file** .
b. From the left pane, expand and select **Network** > **CorpServer** > **BU-Office1** .
c. Select **Save** .
d. Select **Next** .
6. Configure BitLocker encryption.
a. Select **Encrypt entire drive** and then select **Next** .
b. Make sure that **New encryption mode** is selected and then select **Next** .
c. Select **Run BitLocker system check** and then select **Continue** .
d. Select **Restart now** .
The computer will reboot, and the encryption process will run automatically.
e. When the encryption process is complete, select **Close** .
7. Verify that encryption is enabled.
a. From the Windows taskbar, select **File Explorer** .
b. From the left pane, select **This PC** .
c. From the right pane, verify that the **System (C:)** drive shows the *encryption lock* icon.

3.4.9 File Encryption Facts

Encryption of files, directories, and hard drives provides an additional level of data security. File encryption is part of a layered defense strategy and helps to protect confidential data in the event that system data is hacked, lost, or stolen. There are different methods that can be used to encrypt data or entire hard drives.

This lesson covers the following topics:

- Encryption supporting confidentiality
- Disk and file encryption
- Transport encryption and key exchange
- Encrypting File System (EFS)
- PGP and GPG
- BitLocker

Encryption Supporting Confidentiality

If data is encrypted, it does not matter if the disk storing the information is stolen or if it can be intercepted when transferred over a network because the threat actor will not be able to understand or change what has been stolen. This use of encryption fulfills the goal of confidentiality.

When deploying a cryptographic system to protect data assets, consideration must be given to how information could potentially be intercepted. Data can be described as being in one of three states:

- **Data at rest** - is the state when the data is in some sort of persistent storage media.
- **Data in transit (or data in motion)** - is the state when data is transmitted over a network.
- **Data in use (or data in processing)** - is the state when data is present in volatile memory, such as system RAM or CPU registers and cache.

Encrypting megabytes or gigabytes of data is referred to as bulk encryption. Asymmetric encryption and private/public key pairs are not often used for bulk encryption because an asymmetric algorithm cannot process large amounts of data efficiently. The computational overhead is too high when using this type of algorithm to encrypt the contents of a disk or series of network packets.

Therefore, bulk data encryption uses a symmetric cipher, such as AES. A *symmetric cipher* can encrypt and decrypt data files and streams of network traffic quickly. The problem is that distributing a symmetric key is challenging. The more people who know the key value, the weaker the confidentiality property is. Luckily, symmetric keys are only 128 or 256 bits long and can easily be encrypted using a public key. Consequently, most data encryption systems use both symmetric and asymmetric encryption in the following sort of scheme:

- The user generates an asymmetric key pair for the chosen cipher, such as RSA or ECC. The private key portion of this is encrypted so the user must supply their account credential to use it. In this context, the private key is the key encryption key (KEK).
- The system generates a symmetric secret key for the chosen cipher, such as AES256 or AES512. This is referred to as a file or media or data encryption key (DEK). This key is used to encrypt the target data.
- The data encryption key is then encrypted using the public key portion of the KEK.
- To access encrypted data, the user must supply a password or start an authenticated session to use their private key to decrypt the secret key, which can then decrypt the data.

Disk and File Encryption

Data at rest encompasses a great many storage mechanisms. These can be thought of in terms of encryption levels. Lower levels, such as encrypting a whole disk, have the advantage of simplicity but can be complex to manage when multiple users need to access the data. Higher levels, such as applying encryption via the file system or a database management system, can be combined with granular access controls.

Type	Description
Full-disk and partition encryption	Full-disk encryption (FDE) refers to a product that encrypts the whole contents of a storage device, including metadata areas not normally accessible using ordinary OS file explorer tools. FDE also encrypts free space areas. FDE primarily protects against physical theft of the disk. A stolen disk can be mounted on any computer, and the threat actor can take ownership of all the data files. This is not possible if the disk is encrypted because it must be unlocked by the user's credentials to access the decryption key. Many storage devices can perform self-encryption using a cryptographic product built into the disk firmware. A self-encrypting drive (SED) could be a hard disk drive (HDD), solid-state drive

Type	Description
	(SSD), or USB flash drive. The disk firmware implements a cryptoprocessor to store the keys so they are not directly exposed to the OS that mounts the disk. An HDD or SSD can be divided into separate logical areas called partitions. Each partition can be formatted with a different file system and mounted as a drive or volume in the OS. Some disk encryption products might be able to encrypt these partitions selectively rather than the whole disk. The partitions could be encrypted using different keys. For example, a disk could contain boot, system, and data partitions. The boot and system partitions could be left unencrypted as they contain only standard OS files, while the data partition is protected by encryption.
Volume and file encryption	A volume is any storage resource with a single file system. Put another way, a volume is how the OS "sees" a storage resource. The technology underlying the volume might be a removable disk or a partition on an HDD or SSD. It could also be a RAID array. Consequently, a volume encryption product is likely to refer to one that is implemented as a software application rather than by disk firmware. For example, while they might loosely be referred to as "disk encryption," Microsoft's BitLocker and Apple's FileVault products perform volume encryption. A volume encryption product may or may not encrypt free space and/or file metadata. A file encryption product is software that applies encryption to individual files (or perhaps to folders/directories). This might depend on file system support. For example, Microsoft's Encrypting File System (EFS) requires that the volume be formatted with NTFS. Metadata can include a list of files, the file owner, and created/last modified dates. Free or unallocated space can contain data remnants, where a file has been marked as deleted, but the data has not actually been erased from the storage medium. If the device has a TPM or an HSM compatible with the encryption product, the disk/volume/file system can be locked by keys stored in the TPM or HSM.
Database encryption	A structured database stores data in tables. Each table is composed of column fields with a given data type. Records are stored as rows in the table with a value entered for each field. The table data is ultimately stored as files on a volume, but access is designed to be mediated through a database management system (DBMS) running a database language such as Structured Query Language (SQL). Typically, the database is hosted on a server and accessed by client applications. The underlying files could be protected by a disk or volume encryption product running on the server. This will typically have an adverse impact on performance, so encryption is more commonly implemented by the DBMS or by a plug-in. Encryption can be applied at different granular levels. While each DBMS supports different features, the following encryption options, based on Microsoft's SQL Server DBMS, are typical.
Database-level encryption	Database- or page-level encryption and decryption occurs when any data is transferred between disk and memory. This is referred to as transparent data encryption (TDE) in SQL Server. A page is the means by which the database engine returns the data requested by a query from the underlying storage files. This type of encryption means that all the records are encrypted while stored on disk, protecting against theft of the underlying media. It also encrypts logs generated by the database.
Record-level encryption	Many databases contain secrets that should not be known by the database administrator. Public key encryption can solve this problem by storing the private key used to unlock the value of a cell outside of the database.

Type	Description
	Cell/column encryption is applied to one or more fields within a table. This can have less of a performance impact than database-level encryption, but the administrator needs to identify which fields need protection. It can also complicate client access to the data. The encryption/decryption mechanism can work in several ways, but with SQL Server's Always Encrypted feature, the data remains encrypted when loaded into memory. It is only decrypted when the client application supplies the key. The plaintext key is not available to the DBMS, so the database administrator cannot decrypt the data. This allows for the separation of duties between the database administrator and the data owner, which is important for privacy. Some solutions may additionally support record-level encryption. For example, a health insurer's database might store protected health information about its customers. Each customer could be identified by a separate key pair. This key pair would be used to encrypt data at a row/record level. The table contains records separately protected by different keys. This allows fine-grained control over how data can be accessed to meet compliance requirements for security and privacy.

Transport Encryption and Key Exchange

Transport/communication encryption protects data in motion. Various transport encryption products have been developed for different networking solutions. Some examples include the following:

- **Wi-Fi Protected Access (WPA)** secures traffic sent over a wireless network.
- **Internet Protocol Security (IPsec)** secures traffic sent between two endpoints over a public or untrusted transport network. This is referred to as virtual private networking (VPN).
- **Transport Layer Security (TLS)** secures application data, such as web or email data, sent over a public or untrusted network.

As with data-at-rest, an asymmetric cipher is not typically used to encrypt the network data directly because it is too inefficient. Transport encryption products use a system of key exchange . This allows the sender and recipient to exchange a symmetric encryption key securely by using public key cryptography:

- Alice obtains a copy of Bob's RSA or ECC public key, typically via Bob's digital certificate.
- Alice encrypts their message using a secret key cipher, such as AES. In this context, the secret key is referred to as a *session key* .
- Alice encrypts the session key with Bob's public key.
- Alice attaches the encrypted session key to the ciphertext message in a digital envelope and sends it to Bob.
- Bob uses their private key to decrypt the session key.
- Bob uses the session key to decrypt the ciphertext message.

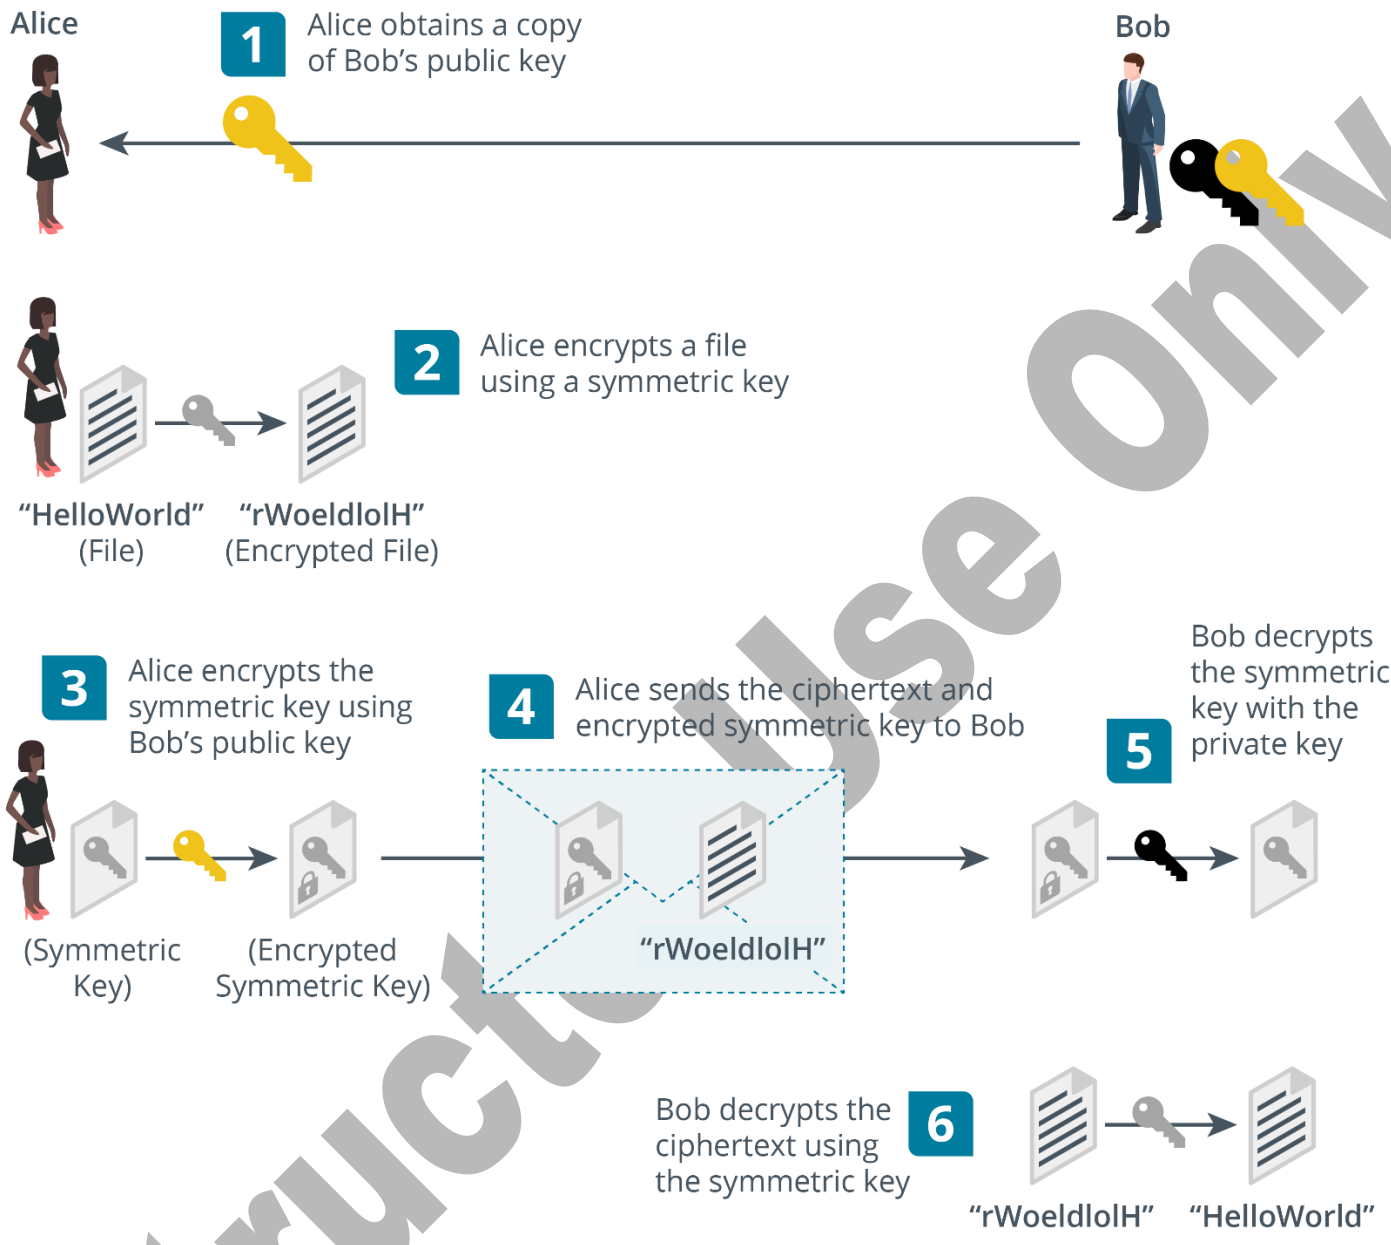

Key exchange using a digital envelope. (Images © 123RF.com.)

Transport encryption also uses cryptography to ensure the integrity and authenticity of messages so that the recipient can verify that they were not modified by someone other than the sender. Integrity and authenticity checking can use a hash-based Message Authentication Code (HMAC) . An HMAC combines the secret key derived during a key exchange with a hash of the message.

Alternatively, the symmetric cipher might be designed to perform Authenticated Encryption (AE). This type of symmetric cipher mode of operation ensures confidentiality and integrity/authenticity.

Encrypting File System (EFS)

The Encrypting File System (EFS) was introduced with NTFS version 3 and has been included in every version of Windows since Windows 2000 except for the Home editions. EFS provides an easy and seamless way for users to encrypt files on their Windows computers. EFS is only used to encrypt individual files and folders.

EFS combines the speed of symmetric encryption with the convenience of asymmetric encryption using a process called key encapsulation. The process for a user to encrypt a file is as follows:

1. The user accesses Properties and, from the General tab, clicks **Advanced** . From there, the user selects Encrypt contents to secure data.
2. Windows generates a pseudo-random number called the file encryption key (FEK). Windows uses the FEK with the AES encryption algorithm to encrypt the file.
3. The FEK is then encrypted using the user's public key. The encrypted FEK is stored in the file's header in a special location called the Data Decryption Field (DDF).
4. The decryption process is the opposite. The user's private key is used first to unlock the DDF and get the FEK. The FEK is then used to decrypt the file.

The encryption and decryption process relies on the user's password being kept safe. If the user account becomes corrupted or the password is forgotten, any encrypted files are lost. To help remedy this, a data recovery agent (DRA) can be set up. The DRA is simply another account that can decrypt the encrypted files.

- The DRA used to be automatically configured as the system administrator in older versions of Windows. Nowadays, it is no longer automatically set up.
- A local DRA for an individual workstation can be configured through the machine's Group Policy settings.
- A domain-wide DRA can be configured in Active Directory. Only a domain administrator can set up a domain-wide DRA.

Additional security considerations are:

- Decryption keys can be backed up to an external USB drive. This ensures access even if the Windows system completely crashes.
- A file is automatically unencrypted when moved or copied to a non-NTFS formatted device or media. A file is also automatically unencrypted when you copy a file over the network using the SMB Protocol.
- Key security relies on the user having a strong password and following proper password security protocols.

PGP and GPG

GNU Privacy Guard (GPG) is an encryption tool that encrypts emails, digitally signs emails, and encrypts documents. GPG is an implementation of the Pretty Good Privacy (PGP) Protocol. PGP is a commercial product now owned by Symantec and makes products that can be used to protect laptops, desktops, USB drives, optical media, and smartphones.

Both PGP and GPG do the following:

- Follow the OpenPGP standard (RFC 4880) for encrypting and decrypting data.
- Combine asymmetric and symmetric cryptography. The process is as follows:
    1. GPG/PGP generates a random symmetric key to encrypt the message.
    2. The symmetric key is then encrypted using the receiver's public key and sent along with the message.
    3. When the recipient receives a message, GPG/PGP first decrypts the symmetric key with the recipient's private key.
    4. The decrypted symmetric key is then used to decrypt the rest of the message.

GPG supports many common algorithms, including RSA, DSA, 3DES, IDEA, MD5, SHA, and more. AES is used by default. PGP can use either RSA or the Diffie-Hellman algorithm for asymmetric encryption and IDEA for symmetric encryption.

BitLocker

Full disk encryption (FDE) means that the entire contents of the drive (or volume), including system files and folders, are encrypted. OS ACL-based security measures are quite simple to circumvent if an adversary can attach the drive to a different host OS. Drive encryption allays this security concern by making the contents of the drive accessible only in combination with the correct encryption key. Disk encryption can be applied to hard disk drives (HDDs) and solid state drives (SSDs).

FDE requires the secure storage of the key used to encrypt the drive contents. Normally, this is stored in a Trusted Platform Module (TPM). The TPM chip has a secure storage area with a disk encryption program, such as Windows BitLocker, to which it can write its keys. It is also possible to use a removable USB drive (if USB is a boot device option). As part of the setup process, you create a recovery password or key. This can be used if the disk is moved to another computer or the TPM is damaged.

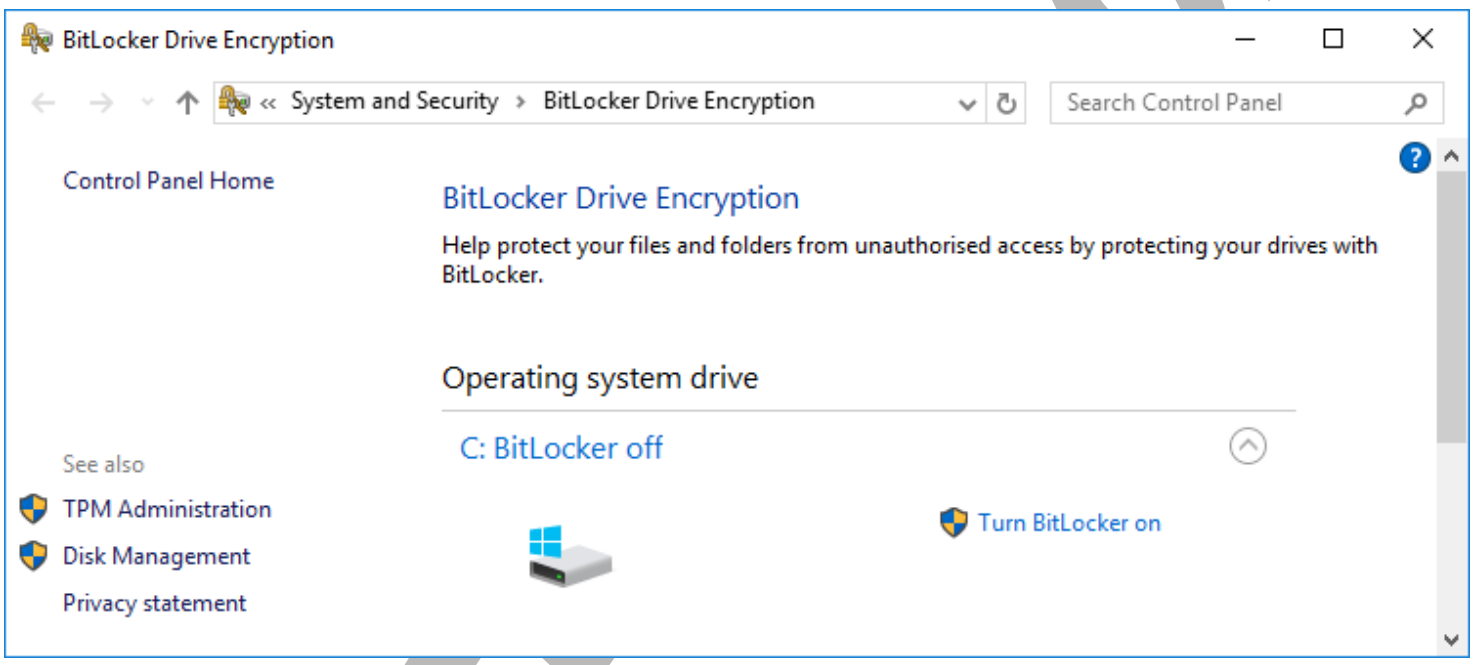

Activating BitLocker drive encryption. (Screenshot used with permission from Microsoft.)

One of the drawbacks of FDE is that because the OS performs the cryptographic operations, performance is reduced. This issue is mitigated by self-encrypting drives (SED) , where the cryptographic operations are performed by the drive controller. The SED uses a symmetric data/media encryption key (DEK/MEK) for bulk encryption. It stores the DEK securely by encrypting it with an asymmetric key pair called the authentication key (AK) or key encryption key (KEK) . The use of the AK is authenticated by the user password. This means that the user password can be changed without having to decrypt and re-encrypt the drive. Early types of SEDs used proprietary mechanisms, but many vendors now develop using the Opal Storage Specification (nvmexpress.org/wp-content/uploads/TCGandNVMe_Joint_White_Paper-TCG_Storage_Opal_and_NVMe_FINAL.pdf), developed by the Trusted Computing Group (TCG).

3.4.10 Practice Questions (Section Quiz)

q_file_encryption_confident_secp8

You create a new document and save it to a hard drive on a file server on your company's network. Then you employ an encryption tool to encrypt the file using AES.

This activity is an example of accomplishing which security goal?

Answers:

- ***Confidentiality**
- Integrity
- Availability
- Non-repudiation

Explanation:

Encrypting a file while it is stored on a hard drive is usually done to provide protection for the object's confidentiality.

Hashing is used to provide integrity.

Using mechanisms like backups and avoiding single points of failure provide availability protection.

Non-repudiation is usually provided for during a secured communication, not while a file is stored on a hard drive.

q_file_encryption_dra_secp8

Which of the following should you set up to ensure encrypted files can still be decrypted if the original user account becomes corrupted?

Answers:

- PGP
- VPN
- ***DRA**
- GPG

Explanation:

If a user account becomes corrupted or the password is forgotten, any encrypted files are lost. To help remedy this, a data recovery agent (DRA) can be set up. The DRA is simply another account that can decrypt the encrypted files.

Pretty Good Privacy (PGP) is an older utility used to encrypt and decrypt data and messages.

A virtual private network (VPN) uses an encryption protocol (such as IPsec, PPTP, or L2TP) to establish a secure communication channel between two hosts or between one site and another site.

GNU Privacy Guard (GPG) is a command line utility that's used to encrypt and decrypt data and messages.

q_file_encryption_encrypt_01_secp8

You want a security solution that protects the entire hard drive and prevents access even if the drive is moved to another system.

Which solution should you choose?

Answers:

- EFS
- VPN

- IPsec
- ***BitLocker**

Explanation:

BitLocker is a Microsoft security solution that encrypts the entire contents of a hard drive, protecting all files on the disk. BitLocker uses a special key that is required to unlock the hard disk. You cannot unlock/decrypt a drive simply by moving it to another computer.

EFS is a Windows file encryption option, but it only encrypts individual files. Encryption and decryption is automatic and dependent upon the file's creator and whether other uses have read permissions.

A virtual private network (VPN) uses an encryption protocol (such as IPsec, PPTP, or L2TP) to establish a secure communication channel between two hosts or between one site and another site. Data that passes through the unsecured network is encrypted and protected.

q_file_encryption_encrypt_02_secp8

Which of the following security solutions would prevent a user from reading a file that they did not create?

Answers:

- ***EFS**
- Bitlocker
- VPN
- IPsec

Explanation:

EFS is a Windows file encryption option that encrypts individual files so that only the user who created the file can open it. Decryption is automatic when the file owner opens it. Other users cannot open the encrypted file unless specifically authorized.

BitLocker is a Microsoft security solution that encrypts the entire contents of a hard drive, protecting all files on the disk. BitLocker uses a special key that is required to unlock the hard disk. You cannot unlock/decrypt a drive simply by moving it to another computer.

A virtual private network (VPN) uses an encryption protocol (such as IPsec, PPTP, or L2TP) to establish a secure communication channel between two hosts or between one site and another site. Data that passes through the unsecured network is encrypted and protected.

q_file_encryption_encrypt_03_secp8

You've used BitLocker to implement full volume encryption on a notebook system. The notebook motherboard does not have a TPM chip, so you've used an external USB flash drive to store the BitLocker startup key.

You use EFS to encrypt the C:\Secrets folder and its contents.

Which of the following is true in this scenario? (Select two.)

Answers:

- ***If the C:\Secrets\confidential.docx file is copied to an external USB flash drive, the file will be saved in an unencrypted state.**
- ***By default, only the user who encrypted the C:\Secrets\confidential.docx file will be able to open it.**
- Any user who is able to boot the computer from the encrypted hard disk will be able to open the C:\Secrets\confidential.docx file.
- The EFS encryption process will fail.
- Only the user who encrypted the C:\Secrets\confidential.docx file is able to boot the computer from the encrypted hard disk.
- If the C:\Secrets\confidential.docx file is copied to an external USB flash drive, the file will remain in an encrypted state.

Explanation:

BitLocker uses full volume encryption, while EFS is used to encrypt individual files and folders. The following are true in this scenario:

- If the C:\Secrets\confidential.docx file is copied to an external USB flash drive, the file will be saved in an unencrypted state.
- Only the user who encrypted the C:\Secrets\confidential.docx file will be able to open it by default.

With BitLocker enabled, any user who has the appropriate startup key or PIN is able to boot the system from the encrypted drive. However, only the user who encrypted the C:\Secrets\ folder will be able to access files within it unless additional user accounts are explicitly added.

q_file_encryption_ipsec_secp8

Which utility would you MOST likely use on OS X to encrypt and decrypt data and messages?

Answers:

- PGP
- ***GPG**
- VPN
- IPsec

Explanation:

GNU Privacy Guard (GPG) is a command line utility that's used to encrypt and decrypt data and messages. GPG is a open source utility and can be used on many different systems, including Windows, Linux, Android, and Apple's OS X.

Pretty Good Privacy (PGP) is an older utility used to encrypt and decrypt data and messages. PGP was purchased a while ago and commercialized. It's owned by NortonLifeLock, formally known as Symantec, and provides products that can protect all sorts of devices, even smartphones. While PGP can be used on OS X, GPG is used by default.

A virtual private network (VPN) uses an encryption protocol (such as IPsec, PPTP, or L2TP) to establish a secure communication channel between two hosts or between one site and another site. A VPN is not used on OS X to encrypt and decrypt data and messages.

IPSec is a protocol used to encrypt VPN communication.

q_file_encryption_pgp_gpg_secp8

Which of the following statements about GPG (GNU Privacy Guard) and PGP (Pretty Good Privacy) is true?

Answers:

- GPG is a proprietary version of PGP.
- GPG and PGP use only symmetric encryption.
- GPG and PGP are incompatible and cannot interoperate.
- ***Both GPG and PGP can be used for encrypting, decrypting, and signing data.**

Explanation:

Both GPG and PGP can be used for encrypting and decrypting data to maintain confidentiality, and for signing data to ensure integrity and authenticity.

GPG is not a proprietary version of PGP. It is an open-source version of PGP.

GPG and PGP use a combination of symmetric and asymmetric encryption, not just symmetric encryption.

GPG and PGP are interoperable. They can work together because they follow the OpenPGP standard.

q_file_encryption_tpm_01_secp8

You would like to implement BitLocker to encrypt data on a hard disk, even if it is moved to another system. You want the system to boot automatically without providing a startup key on an external USB device.

What should you do?

Answers:

- ***Enable the TPM in the BIOS.**
- Disable USB devices in the BIOS.
- Use a PIN instead of a startup key.
- Save the startup key to the boot partition.

Explanation:

When a system boots, the startup key is required to unlock the encrypted volume. The system startup key can be saved in the Trusted Platform Module (TPM). With the startup key saved in the TPM, the system can start without additional intervention.

The system will not start without the startup key. Without a TPM, the startup key must be stored on a USB drive.

You can require a PIN in addition to a startup key, but the PIN cannot replace the startup key.

Storing the startup key on the boot drive would expose it to compromise.

q_file_encryption_tpm_02_secp8

You want to protect data on hard drives for users with laptops. You want the drive to be encrypted, and you want to prevent the laptops from booting unless a special USB drive is inserted. In addition, the system should not boot if a change is detected in any of the boot files.

What should you do?

Answers:

- Implement BitLocker without a TPM.
- ***Implement BitLocker with a TPM.**
- Have each user encrypt user files with EFS.
- Have each user encrypt the entire volume with EFS.

Explanation:

If you use BitLocker without a TPM, system integrity checks are not performed. The TPM is required for saving the startup file information that is used to verify system integrity. When using BitLocker without a TPM, you must use a startup key on a USB device. When using a TPM, this is an optional configuration.

Use BitLocker to encrypt the entire system volume and protect both operating system and user data. Use BitLocker with a Trusted Platform Module (TPM) to protect the boot environment components such as the BIOS, Master Boot Record, Boot Sector, Boot Manager, and Windows Loader. The system is shut down if a boot environment change is detected. Using BitLocker, drives are locked if they are moved to another computer, and you can require a startup key on a USB drive or a PIN before the system boots.

EFS encrypts individual files. With EFS, only the user who encrypted the file and any additionally designated users can access the file. EFS does not provide integrity checks for boot files.

q_file_encryption_trans_data_secp8

Which of the following database encryption methods encrypts the entire database and all backups?

Answers:

- ***Transparent Data Encryption (TDE)**
- Column-level
- Application-level
- Bitlocker

Explanation:

Transparent Data Encryption (TDE) encrypts the entire database and all backups. TDE:

- encrypts data at rest, which is data not being currently used.
- is called transparent because when an authorized user needs to access the data, it is automatically decrypted so the user does not see the process or need to do anything extra.

Column-level encryption allows the administrator to encrypt each column separately.

In application-level encryption, the program that is used to create or modify the data is responsible for encrypting the data.

BitLocker is a Microsoft security solution that encrypts the entire contents of a hard drive, protecting all files on the disk.

q_file_encryption_unencrypt_secp8

You have transferred an encrypted file across a network using the Server Message Block (SMB) Protocol.

What happens to the file's encryption?

Answers:

- The encryption carries over to the new location.
- The encryption inherits from the new location.
- An encrypted file cannot be moved using SMB.
- ***The file is unencrypted when moved.**

Explanation:

A file is automatically unencrypted when you copy it over a network using the SMB Protocol.

The encryption does not carry over to the new location, nor does the file inherit from the new location.

A file can be moved using the SMB Protocol.

3.5 Public Key Infrastructure

As you study this section, answer the following questions:

- What is the lifecycle of an encryption key?
- What is the role of a certificate authority (CA)?
- What are the types of certificates?
- Which standard defines the format of certificates?
- Which trust model would be used to connect the CAs of two organization's?

In this section, you will learn to:

- Manage certificates.

The key terms for this section include:

Term	Definition
Public key infrastructure (PKI)	A framework of certificate authorities, digital certificates, software, services, and other cryptographic components deployed for the purpose of validating subject identities.
Third party CAs	In PKI, a public CA that issues certificates for multiple domains and is widely trusted as a root trust by operating systems and browsers.
Digital certificate	Identification and authentication information presented in the X.509 format and issued by a certificate authority (CA) as a guarantee that a key pair (as identified by the public key embedded in the certificate) is valid for a particular subject (user or host).
Public Key Cryptography Standards (PKCS)	A series of standards defining the use of certificate authorities and digital certificates.

Certificate signing request (CSR)	A Base64 ASCII file that a subject sends to a CA to get a certificate.
Common name (CN)	An X500 attribute expressing a host or username, also used as the subject identifier for a digital certificate.
Subject alternative name (SAN)	A field in a digital certificate allowing a host to be identified by multiple host names/subdomains.
Wildcard	In PKI, a digital certificate that will match multiple subdomains of a parent domain.
Certificate revocation list (CRL)	A list of certificates that were revoked before their expiration date.
Online Certificate Status Protocol (OCSP)	Allows clients to request the status of a digital certificate, to check whether it is revoked.
Root certificate	In PKI, a certificate authority that issues certificates to intermediate certificate authorities in a hierarchical structure.
Certificate chaining/Chain of trust	A method of validating a certificate by tracing each CA that signs the certificate up through the hierarchy to the root CA.
Self-signed certificate	A digital certificate that has been signed by the entity that issued it, rather than by a certificate authority.
Escrow	In key management, the storage of a backup key with a third party.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions.
	• Public key infrastructure (PKI)
	  ○ Public key
	  ○ Private key
	  ○ Key escrow
	• Digital signatures
	• Certificates
	  ○ Certificate authorities

	○ Certificate revocation lists (CRLs) ○ Online Certificate Status Protocol (OCSP) ○ Self-signed ○ Third-party ○ Root of trust ○ Certificate signing request (CSR) generation ○ Wildcard 4.1 Given a scenario, apply common security techniques to computing resources. • Application security ○ Input validation ○ Secure cookies ○ Static code analysis ○ Code signing
TestOut Security Pro	4.2 Implement Encryption Technologies 4.2.3 Manage Certificates

3.5.1 Public Key Infrastructure (Lesson Video)

Transcript:

When we encrypt data over the internet, we generally utilize asymmetric encryption methods that involve sending a user's public key to provide confidentiality and trust. Because these keys are public, we need a way to manage and protect them.

Key management covers these keys' whole life cycle. During this cycle, we must keep keys safe because we need to be sure that the public key we're using really does belong to the organization it's associated with. Public key infrastructure, or PKI, handles this for us.
A PKI provides an environment where public encryption keys can be created and managed. At the heart of a PKI are Certificate Authorities, which are responsible for issuing, validating, and revoking certificates. In this lesson, I'll go over the concept of Certificate Authorities, or CAs, and the process they use to verify a certificate with all of its different attributes.
A PKI relies on certificates that validate organizations. This creates a web of trust across the internet, allowing us to perform transactions confidently with websites around the globe. A PKI requires several elements to be effective.
The first element is the CA. CAs need to be reputable organizations that are respected enough to issue public certificates to organizations that want to communicate securely over the internet.
To increase security, CAs operate in a hierarchy of multiple CAs. This is done so that if one CA is compromised, only the certificates it issued need to be revalidated.
The first CA is the Root CA. This is a self-signed certificate that's used to validate additional CAs.
These Subordinate CAs are also known as Intermediate CAs. We can have multiple Intermediate CAs based on their policies and regulations. The Intermediate CAs validate the Issuing certificate authorities, and the Issuing CA is the one that hands out the certificates.
Now that we understand the certificate authorities' hierarchy, let's look at how an organization can obtain their very own certificate.
To obtain a certificate, an organization needs to first send in a certificate signing request, or CSR, to a Certificate Authority. The CSR should contain the organization's public key, domain name, and digital signature. Then the CA verifies this information and issues the certificate.

When filling out the CSR, the organization provides their Common Name, or CN, which is more commonly referred to as the Fully Qualified Domain Name. For example, here you see that TestOut's Common Name would be www.testout.com.
The organization can also apply for a Subject Alternative Name, or SAN. This allows one certificate to apply to multiple host names. For example, TestOut could apply for a SAN that would cover site1.testout.com and site2.testout.com.
Once the CA has received the CSR, they verify the information and provide the certificate to the requesting organization.
Sometimes the CA relies on a third party to perform the validation. These third parties are called Registration Authorities. An RA is certified by a Root CA and is authorized to issue certificates for specific uses only. No matter who issues the certificates, each one has specific attributes for specific purposes.
Each CA's responsibility is to maintain a database that contains information on each certificate they've issued. This information is mainly their certificates' attributes.
These attributes contain everything from the serial number and signature algorithm to the public key and expiration date. CAs use the X.509 standard to define these attributes.
Each certificate has an expiration date. Before the certificate expires, the organization must revalidate the information and renew their certificate. If they don't, it's no longer valid.
Aside from expiration, there are other reasons a CA might invalidate or revoke a certificate. For example, if the organization is found to no longer exist, if the private key is compromised, or if the certificate was discovered to be fake, the CA should immediately revoke the certificate.
When a CA does revoke a certificate, it's added to a Certificate Revocation List, or CRL. This is a type of certificate blacklist. CAs maintain the CRL as part of their databases, and these should be updated quite often.
The X.509 standard also defines an internet protocol that can be used to determine a certificate's current state. This is called the Online Certificate Status Protocol, or OCSP.
The way OCSP works is your browser sends a status request to an OCSP responder and receives a response as to whether a certificate is valid or has been revoked.
Using the OCSP provides a few benefits. These include providing more timely information on a certificate's status, better bandwidth management because the client doesn't need to download the CRL, and a grace period for expired certificates.
If we follow the guidelines set in the X.509 standard, we can be assured that internet certificates are valid and have been checked appropriately by reputable Certificate Authorities.
That's it for this lesson. In this lesson, we covered Certificate Authorities and the hierarchical structure they're set up in. We also looked at the process an organization goes through to request a certificate for themselves. Finally, we looked at some certificate attributes and how CAs use them to validate certificates or put them on a Certificate Revocation List.

3.5.2 Public Key Infrastructure Facts

Public Key Infrastructure (PKI) is the framework that helps to establish trust in the use of public key cryptography to sign and encrypt messages via digital certificates. A digital certificate is a public assertion of identity validated by a certificate authority (CA).

This lesson covers the following topics:

- Certificate authorities
- Digital certificates
- Certificate attributes

Certificate Authorities

Public key cryptography solves the problem of distributing encryption keys when you want to communicate securely with others or authenticate a message that you send to others.

- When you want others to send you confidential messages, you give them your public key to encrypt the message. The message can then only be decrypted by your private key, which you keep known only to yourself.

- When you want to authenticate yourself to others, you sign a hash of your message with your private key. You give others your public key to use to verify the signature. As only you know the private key, everyone can be assured that only you could have created the signature.

The basic problem with public key cryptography is that while the owner of a private key can authenticate messages, there is no mechanism for establishing the owner's identity. This problem is particularly evident with e-commerce. How can you be sure that a shopping site or banking service is really maintained by whom it claims? The fact that the site is distributing a public key to secure communications is no guarantee of actual identity. How do you know that you are corresponding directly with the site using its genuine public key? How can you be sure there is no threat actor with network access intercepting and modifying what you think the legitimate server is sending you?

Public key infrastructure (PKI) aims to prove that the owners of public keys are who they say they are. Under PKI, anyone issuing a public key should publish it in a digital certificate. The certificate's validity is guaranteed by a certificate authority (CA).

PKI can use private or third-party CAs . A private CA can be set up within an organization for internal communications. The certificates it issues will only be trusted within the organization. For public or business-to-business communications, a third-party CA can be used to establish a trust relationship between servers and clients. Examples of third-party CAs include Comodo, DigiCert, GeoTrust, IdenTrust, and Let's Encrypt.

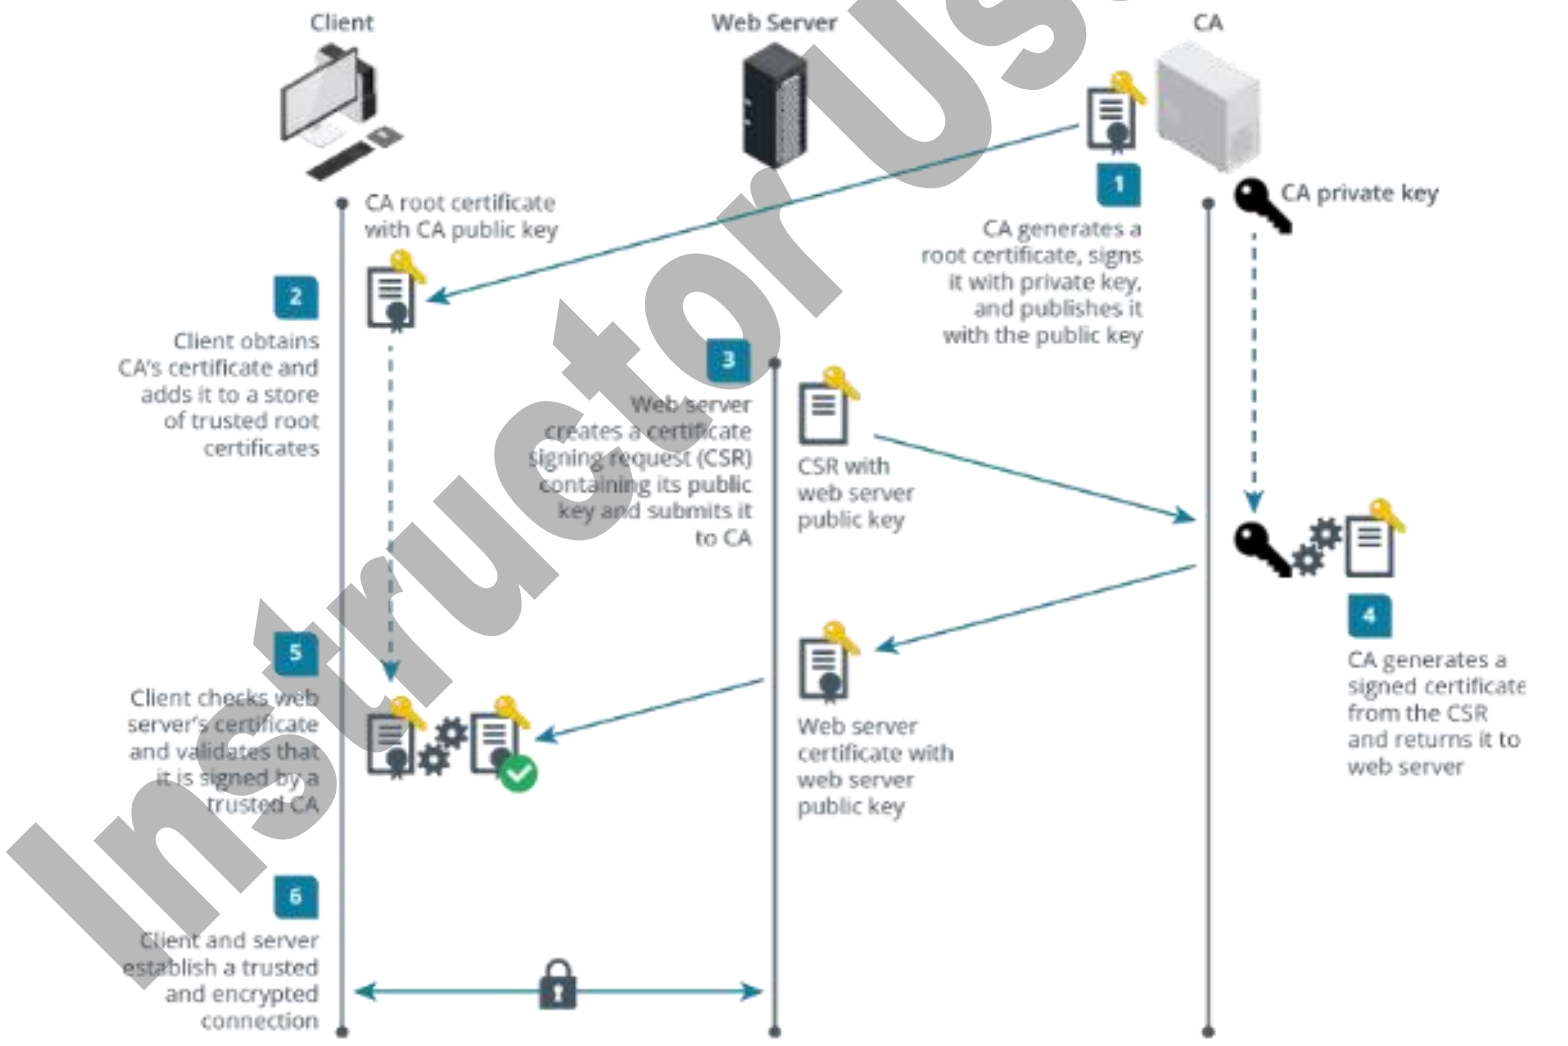

Public key infrastructure allows clients to establish a trust relationship with servers via certificate authorities.

The functions of a third-party public CA are as follows:

- Provide a range of certificate services useful to the community of users serviced by the CA.
- Ensure the validity of certificates and the identity of those applying for them (registration).
- Establish trust in the CA with users, governments, regulatory authorities, and enterprises such as financial institutions.
- Manage the servers (repositories) that store and administer the certificates.
- Perform key and certificate lifecycle management, notably revoking invalid certificates.

Digital Certificates

A digital certificate is essentially a wrapper for a subject's public key. As well as the public key, it contains information about the subject and the certificate's issuer. The certificate is digitally signed to prove that it was issued to the subject by a particular CA. The subject could be a human user (for certificates allowing the signing of messages, for instance) or a computer server (for a web server hosting confidential transactions, for instance).

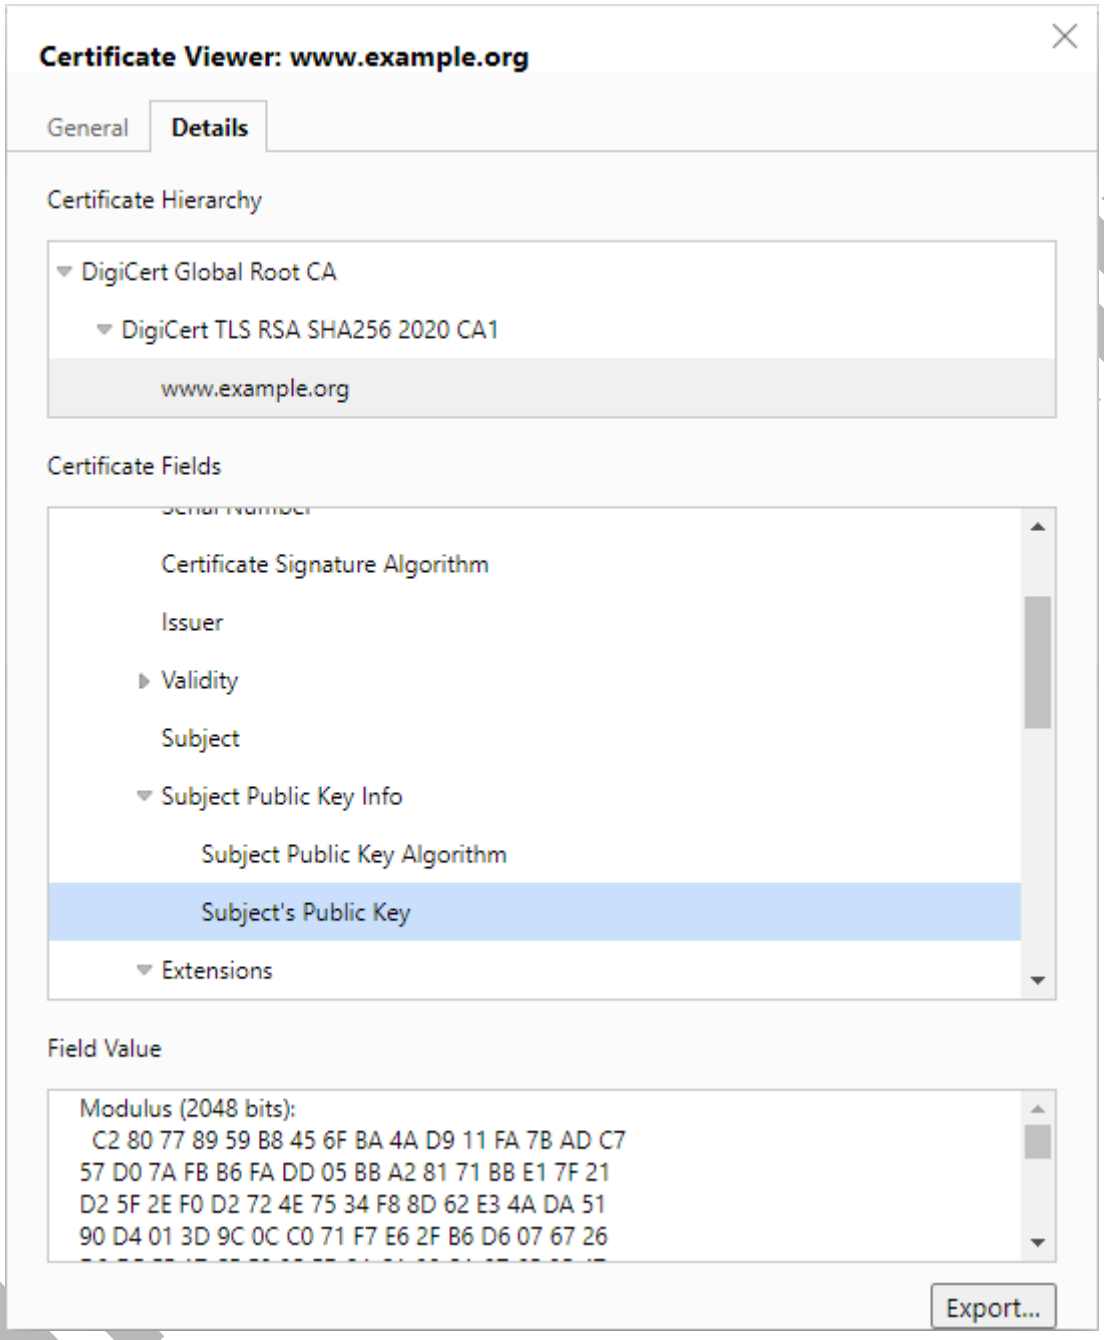

Digital certificate details showing the subject's public key. (Screenshot used with permission from Microsoft.)

Digital certificates are based on the X.509 standard approved by the International Telecommunications Union and standardized by the Internet Engineering Task Force (tools.ietf.org/html/rfc5280). RSA also created a set of standards, referred to as Public Key Cryptography Standards (PKCS) , to promote the use of public key infrastructure.

Certificate Attributes

Each CA is responsible for maintaining a database containing the information or attributes of each certificate. The attributes that can be included are:

- Version - The X.509 version used for the certificate.
- Serial Number - A unique identifier for each certificate.
- Signature algorithm - The algorithm used to sign the certificate (SHA-2, RSA, etc.).
- Issuer - The CA that issues the certificate.
- Valid From and Valid To - The two fields that show the validity period of the certificate.
- Subject - The field that contains the name and location information of the organization.
- Public Key - The algorithm used to create the key and the public key information.

Depending on the organization, there may also be additional optional fields called extensions in the certificate.

One of the key attributes is the Valid To field. If a certificate is not renewed by this date, it will expire and no longer be valid. Aside from expiration, some other reasons a certificate might be invalidated are:

- The organization no longer exists.
- The private key has been compromised.
- The issued certificate is discovered to be fake.

If a certificate is invalidated for these or other reasons, it will be added to a certificate revocation list (CRL). The CRL is a blacklist of certificates. CAs must maintain and constantly update the CRLs as part of their databases. Web browsers automatically download updated CRLs at set intervals.

The X.509 standard also defines an internet protocol that can be used to determine the validity or state of a certificate. This is called the Online Certificate Status Protocol (OCSP). OCSP can be used to simplify the process of checking whether or not a certificate is valid.

OCSP is designed to replace CRLs. Instead of a CA maintaining the CRL, an OCSP server called a responder maintains the lists of any revoked certificate. When the browser connects to a site, the browser sends a request to the OCSP responder to check the validity of the certificate. OCSP provides the following benefits:

- Timely information on the status of a certificate.
- Better bandwidth management because the client does not download the entire CRL.
- A grace period for expired certificates.

3.5.3 Certificate Types (Lesson Video)

Transcript:

Depending on your purpose, there are different types of public key infrastructure certificates, or PKI certificates. These certificates are used to verify an organization's identity and ownership of the public key. In this lesson, I'll cover these types of certificates and how you can use them. I'll also go over some different validation levels used for SSL certificates.

A root certificate is the first certificate that a Certificate Authority, or CA, creates. This certificate is self-signed and is used to sign lower-level certificates, such as intermediate certificates. These certificates go through different processes to be approved, and the process depends on the certificate and organization.
Subject Alternative Name certificates, or SAN certificates, allow an organization to cover multiple domains. For example, TestOut could cover testout.com, testout.net, or even labsim.com with the same certificate.
Wildcard certificates are like SAN certificates. Wildcard certificates allow an organization to cover unlimited subdomains. For example, TestOut could cover any site that ended in testout.com with the same wildcard certificate.
Code-signing certificates are used by app developers to prove that an application is legitimate and hasn't been altered or compromised. When the user installs a program and there's no certificate, they receive a warning that they shouldn't trust the app.

Self-signed certificates are certificates that haven't been validated and signed by a CA. These certificates are easy and free to make, but they don't provide the same protection and security as a CA-validated certificate. If a website uses a self-signed certificate, the user sees a warning when visiting the site.
Secure, encrypted emails are generally sent using the S/MIME Protocol. When you send secure emails, you need to know the recipient's public key, which you find in email certificates. These certificates are usually used within an organization that uses its own CA, but some public CAs also provide email certificates.
You can use certificates to identify and validate individual users and computers. You generally see these kinds of certificates in a network environment where they're used to validate a computer or user to the server.
The most common use for certificates is for websites that use SSL or TLS. These certificates validate that the website is authentic and secure so that users feel safe.
When a site wants to purchase a certificate, there are three different validation levels a CA can offer. These are domain validation, organization validation, and extended validation.
Domain validation is the lowest level and isn't very secure. A CA issues a domain-validated certificate to anyone who's listed as the domain administrator in the WHOIS record. The CA usually validates this with a simple phone call or email, which isn't hard to spoof. These certificates can be issued usually within a matter of minutes.
An organization validation requires the certificate purchaser to not only prove that they're a domain administrator, but also that the organization is real. This process varies with each CA, but it's more in-depth than a domain validation. These certificates may take 1 to 3 days to be issued.
Extended validation is the highest level of trust. An extended validation requires the purchaser to prove that they're the domain administrator and also requires a much more thorough vetting of the organization itself. These certificates may take 1 to 5 days to be issued.
That's it for this lesson. In this lesson, we discussed the different types of PKI certificates and how CAs implement them. We also looked at the three validation levels for SSL certificates and how each level differs.

3.5.4 Certificate Types Facts

PKI certificates are used to verify an organization’s identity and ownership of a public key. When an organization requests a certificate, they must choose which type they need. The Certificate Authority needs to validate the organization before issuing the certificate. The level of validation depends on the certificate type being requested.

This lesson covers the following topics:

- Certificate types
- Certificate signing requests

Certificate Types

Depending on the use and situation, there are different types of public key infrastructure (PKI) certificates. The following table explains what these certificate types are and how they can be used:

Certificate Type	Description
Root certificate	A root certificate is the first certificate that a Certificate Authority creates. Root certificates are: • Self-signed certificates. These certificates go through a different validation process, which varies depending on the certificate and organization. • Used to sign lower-level certificates such as intermediate certificates.
Subject Alternative Name (SAN) certificate	SAN certificates allow an organization to cover multiple domains with one certificate. For example, TestOut could cover the following domains in a single SAN certificate:

Certificate Type	Description
	• TestOut.com • TestOut.net • LabSim.com
Wildcard certificate	Wildcard certificates are similar to SAN certificates. But instead of covering multiple domains, the organization can cover one domain and multiple subdomains. For example, TestOut could cover the following in one certificate: • quiz.testout.com • labs.testout.com • videos.testout.com
Code-signing certificate	Code-signing certificates are used by app developers to prove their application is legitimate. If a user tries to run an app that does not have a certificate, they will receive an error stating that the app cannot be trusted. The user can decide to close the app or run it.
Self-signed certificate	Self-signed certificates are certificates that have not been validated or signed by a CA. • Self-signed certificates are easy and free to make. • Self-signed certificates do not provide the same protection and security as a CA-validated certificate. • When a user visits a website using a self-signed certificate, they see a warning that the certificate is not trusted.
Email certificate	Secure, encrypted emails are sent using the S/MIME Protocol. • Senders need to know the recipient's public key when sending a secure email. The public key is found in email certificates. • Email certificates are mainly used in an organization that uses its own CA. However, some public CAs provide email certificates as well.
User and computer certificate	User and computer certificates are used in a network environment to identify and validate specific users or computers. When a user or computer logs into a network, their certificate is sent to the server for validation. This provides extra security to the network.

Certificate Signing Requests

Registration is the process by which end users create an account with the CA and become authorized to request certificates. The exact processes by which users are authorized and their identity proven are determined by the CA implementation. For example, in a Windows domain network, users and devices can often auto-enroll with the CA just by authenticating to Active Directory. A third-party CA might perform a range of tests to ensure that a subject is who they claim to be. It is in the CA's interest to ensure that it only issues certificates to legitimate users, or its reputation will suffer.

When a subject wants to obtain a certificate, it first generates a key pair comprising private and public asymmetric keys for the chosen cipher, such as RSA or ECC, and key length. The private key must be kept well protected and known only to the subject. The subject then completes a certificate signing request (CSR) and submits it to the CA. The CSR is a file containing the information the subject wants to use in the certificate, including its public key.

The CA reviews the certificate and checks that the information is valid. For a web server, this may simply mean verifying that the subject name and fully qualified domain name (FQDN) are identical and verifying that the CSR was initiated by the person administratively responsible for the domain, as identified in the domain's WHOIS records. If the request is accepted, the CA signs the certificate and sends it to the subject.

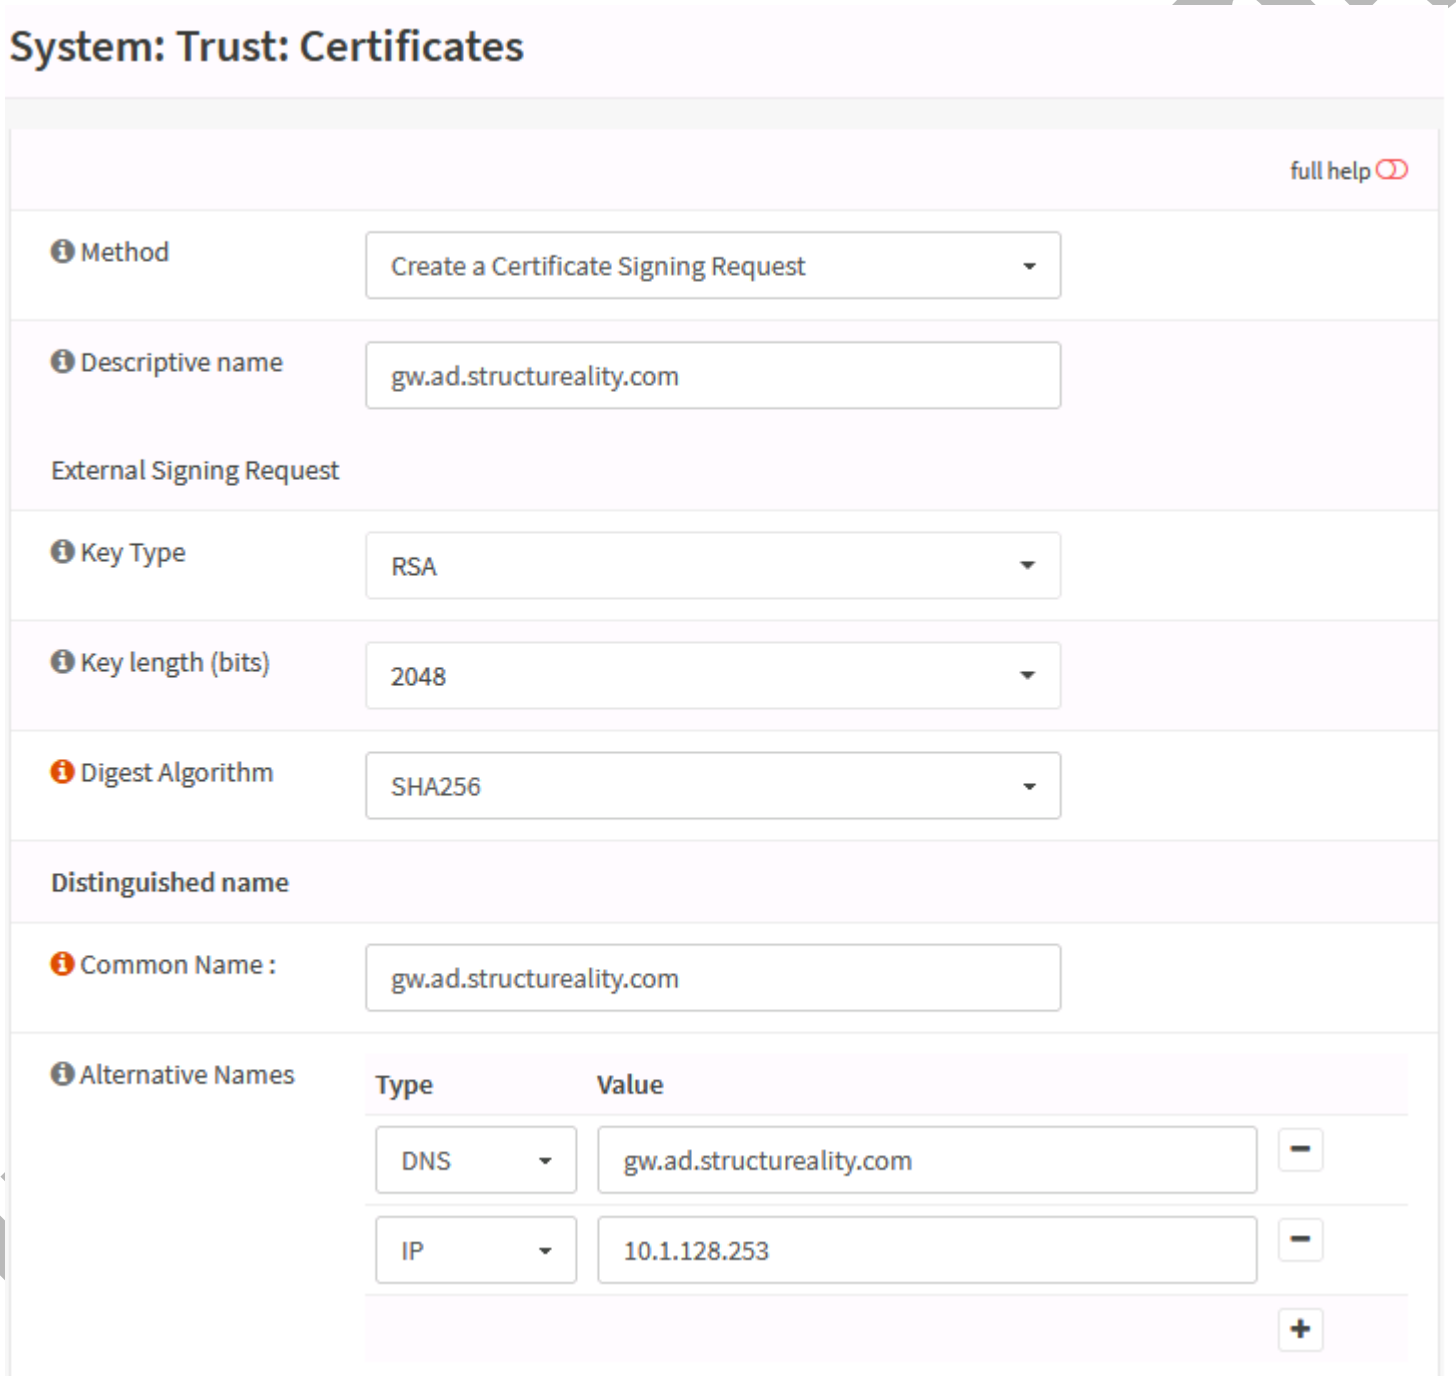

Using a web form in the OPNsense firewall appliance to request a certificate. The DNS and IP alternative names must match the values that clients will use to browse the site.

When certificates were first introduced, the common name (CN) attribute was used to identify the fully qualified domain name (FQDN) by which the server is accessed, such as www.comptia.org . This usage grew by custom rather than design, however. The CN attribute can contain different kinds of information, making it difficult for a browser to interpret it correctly. Consequently, the CN attribute is now deprecated as a method of validating a subject identity that needs to resolve to some type of network address.

The subject alternative name (SAN) extension field is structured to represent different types of identifiers, including FQDNs and IP addresses. If a certificate is configured with a SAN, the browser should validate that and ignore the CN value.

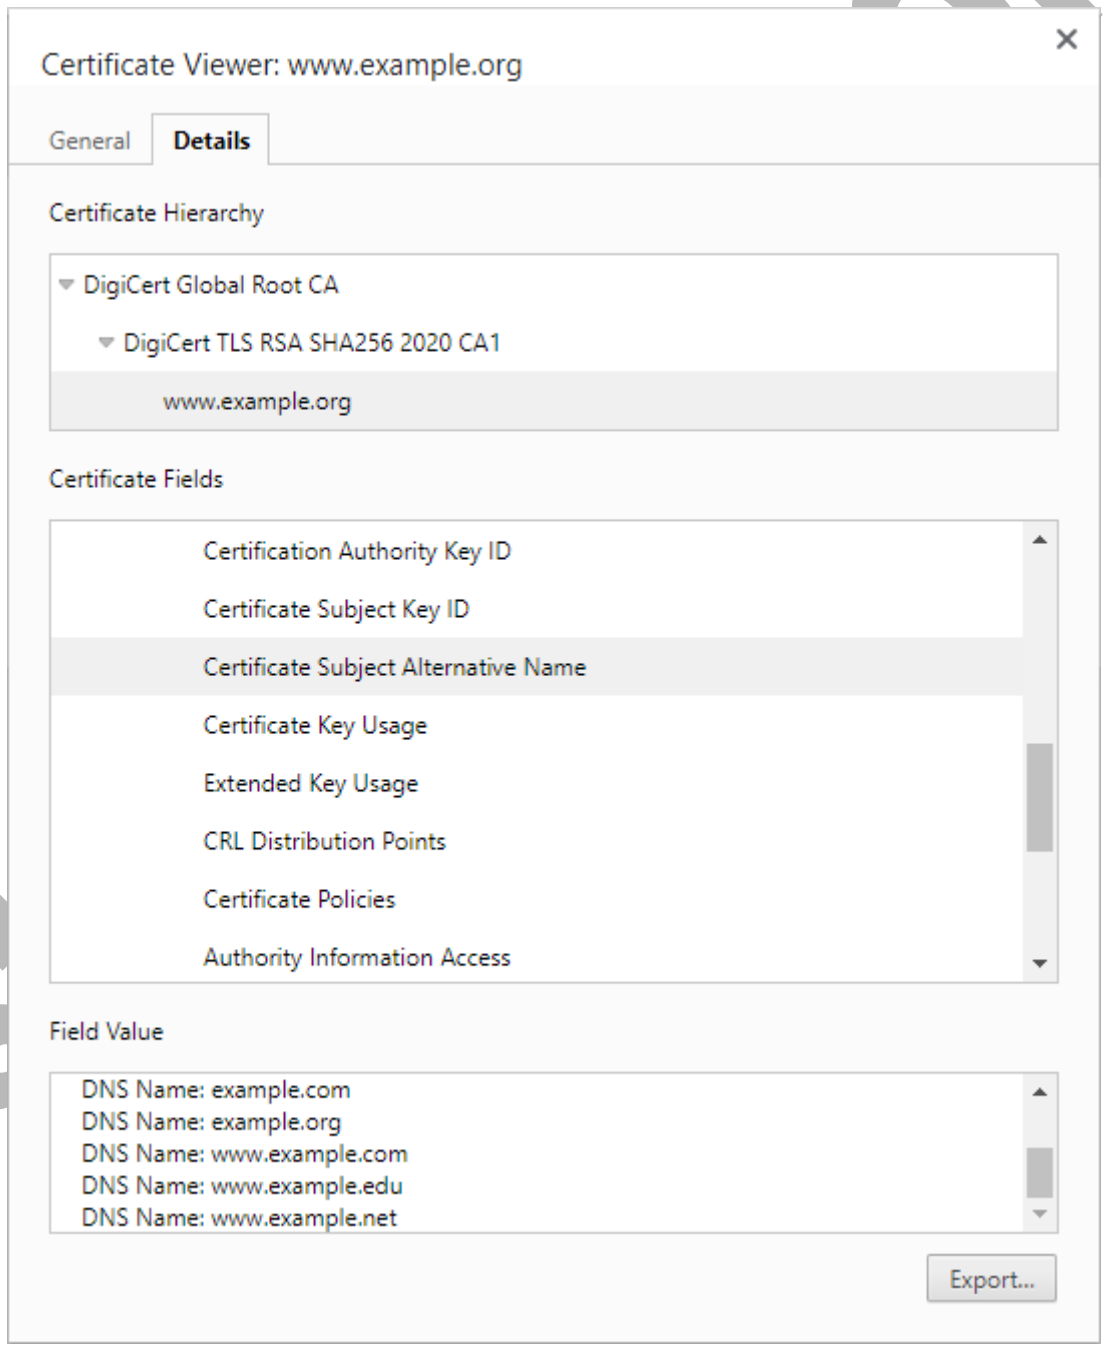

The example domain's certificate is configured with alternative subject names for different top-level domains and subdomains. (Screenshot used with permission from Microsoft.)

It is still safer to put the FQDN in the CN because not all browsers and implementations stay up to date with the standards.

The SAN field also allows a certificate to represent different subdomains, such as www.comptia.org and members.comptia.org . Listing the specific subdomains is more secure, but if a new subdomain is added, a new certificate must be issued. A **wildcard** domain, such as *.comptia.org , means that the certificate issued to the parent domain will be accepted as valid for all subdomains (to a single level).

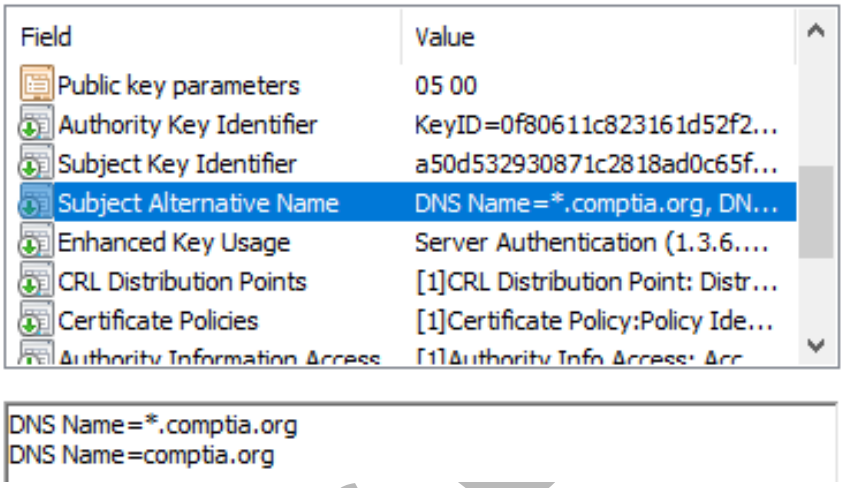

CompTIA's website certificate configured with a wildcard domain, allowing access via either https://comptia.org or https://www.comptia.org. (Screenshot used with permission from Microsoft.)

A certificate also contains fields for Organization (O), Organizational Unit (OU), Locality (L), State (ST), and Country (C). These are concatenated with the common name to form a Distinguished Name (DN). For example, Example LLC's DN could be: CN=www.example.com, OU=Web Hosting, O=Example LLC, L=Chicago, ST=Illinois, C=US .

Different certificate types can be used for purposes other than server/computer identification. User accounts can be issued with email certificates, in which case the SAN is an RFC 822 email address. A code-signing certificate is used to verify the publisher or developer of software and scripts. These do not use a SAN, but the CA must validate the organization and locale details to ensure accuracy and that a rogue developer is not attempting to impersonate a well-known software company.

3.5.5 Manage Certificates (Demo Video)

Transcript:

In this demonstration, we're going to practice managing the public key infrastructure.

To do this, we'll use the Windows Server 2022 system that already has Active Directory Certificate Services installed, which makes it a certificate authority, or CA.
It's called Active Directory Certificate Services because the CA itself is integrated with Active Directory to authenticate certificate requests.
We're going to practice requesting, approving, issuing, and revoking certificates on the CA. In order to manage the CA, we need to launch the CA console by clicking Tools and Certificate Authority.
We can see here the CA itself and the certificate categories: Revoke Certificates, Issue Certificates, Pending Requests, and Failed Requests.
By default, the CA is set to automatically approve any new certificate requests. However, we don't want that to happen.

Instead, we want to move the certificate request to a pending state, here in Pending Requests, until an administrator comes through and manually reviews the request. Then he or she either approves or denies it.
To make sure that this is set correctly, we're going to right-click on the CA, go to Properties, go to Policy Module, go to Properties again. We'll change this so that new certificate requests are sent to Pending until approved.
Now we are given a prompt that we need to restart certificate services for that change to go into effect. I'll do that now by right clicking on the CA, going to All Tasks, then Stop Service. Then I'll do the same thing and Start the service again.
With that setting checked, we want to request a new certificate now. There are two different ways to do this.
With the first method, I install the certificate authority and I also install the Certificate Enrollment Page. This allows users to request certificates using a web page running from the Internet Information Server Service on the CA itself.
Let's go ahead and open up Microsoft Edge. We're going to point it to the right URL for the local host. There it is. It's going to be local host /certsrv. You see here the first task on the list is request a certificate.
Alright, we're going to come back to this window a little bit later. The second method is using the MMC console. Let's go ahead and right-click on the window. Go to Run, MMC, and hit Yes.
First, we're going to need to add some snap-ins. So, we're going to go Add or Remove Snap-in. We need the Certificate snap-in. Go to Add, My User Account, Finish, and Okay. The snap-in is here.
Now we can see all the different types of certificates that we can create for our current user. We're going to actually create one under Personal.
Go to Certificates, right-click, click All Tasks, and Request New Certificate. Go through Next, Next, Select User, Enroll, and Finish.
Now let's go back to our CA console. If we go into Pending Requests, we can see that it was properly sent straight to our Pending Requests. It wasn't actually issued.
From this screen, we can view the pending request and see the request ID here.
This is where we can now request whether we want to approve or deny this certificate. If we right-click it, we have the option to issue or deny. Let's go ahead and issue this certificate.
Now the certificate request is gone from Pending. If we go to Issued Certificates, we can now see It's been officially issued.
We can see all sorts of information here. For example, who it's issued to, who it was issued by, and the certificate authority, which is our certificate authority.
We see the certification expiration date here and its intended purpose. So, that's a personal certificate. It's intended for EFS encryptions, securing email, and client authentication. It was created using the user certificate template.
If we double-click the certificate, we can view the date range when it's valid and we can go to the Details tab. Here we can view additional details about the certificate. And if we click the Certificate Path tab, we can view the certificate path up to the root CA.
In this demonstration, we only have one CA and there are no subordinate CAs. So, the certificate path is actually really short. But in many organizations, you'll have one or more root CAs with subordinate CAs associated with it.
So, the Certificate Path tab can be used to view the path through the subordinate CAs up to the root CA. Let's close this and go back to our CA console.
There may be different situations when the certificate shouldn't be used. For example, this is the case if the private key gets compromised in some way. In these situations, we should no longer use this certificate for encryption.
We would need to revoke it. To do this, we simply right-click the certificate, go to All Tasks, and to Revoke. You can specify a reason code for why this certificate is being revoked.
If you look, the last code is Certificate Hold. This is used in situations where you think there might be a problem with the certificate but you're not 100 percent sure.
We want to put the certificate on hold without fully revoking it so that we can verify whether or not it's an issue that would require it to be fully revoked or not. If there is no issue, we can take it off hold and continue using it. If there's a problem, we can fully revoke it.
If you were to use one of the other options here to revoke the certificate, such as Key Compromise, you couldn't unrevoke it. It's gone. You'll have to issue a new certificate. Let's go ahead and put it in Certificate Hold and hit Yes.
You see it clears out the Issued Certificates, and now it's on hold. Let's go to Revoke. You can see right here the revocation reason: Certificate Hold.
Because we just held it instead of revoking it all the way, we can unrevoke it. To do this, we right-click, go to All Tasks, and Unrevoke Certificate. Alright, it removed it.
If we go back to Issued Certificates, we can see the certificate is back for use again. Let's go ahead and do this again. But this time, go to All Tasks and revoke it for a different reason.
Let's go ahead and just say Key Compromise. Hit Yes. Alright, now it's revoked. If we go back to Revoked Certificates and we right-click, go to All Tasks, and try to unrevoke it, we can't.

We get this error: unrevoke command failed. Certificates can only be unrevoked if they're revoked with reason code Certificate Hold.
So, this certificate no longer can be used. You would have to reissue a new certificate in order for them to use it.
After you've revoked a certificate, you need to let everyone know that the certificate has been revoked and shouldn't be used anymore. This is done by publishing the latest Certificate Revocation List.
To do this, we go to our Revoked Certificates folder, we right-click, go to All Tasks, and then click Publish. This will publish the latest list of revoked certificates to the Certificate Revocation List.
There are two different types of Certificate Revocation Lists, or CRLs. You can publish a full CRL, which contains a list of all the revoked certificates. Or we can just publish a delta CRL which only contains the certificates that have been revoked since the last time the full CRL was published.
For our purposes, we're going to publish a new CRL because it's the first time we've actually published a CRL from this CA. Let's go ahead and click Okay.
If you forget to do this manually, it's okay because the CA should be configured to automatically publish the CRL at a particular interval. If we right-click, hit Revoke Certificates again, and go to Properties this time, we can see the CRL publishing parameters.
By default, the full CRL is published once a week and the delta CRLs are published once a day.
If we click the View CRLs tab, then click CRL, and then click the revocation list, we can see the current CRL and the certificates that were just revoked. Because this CA hasn't been in use very long, there are not many certificates to be revoked. Let's go ahead and close it. Click OK and OK.
Alright, now let's look at the properties of the CA. Let's right-click the Root CA and Properties. Alright, we're going to go to the Extensions tab. The Extensions tab has two acronyms that are very important to understand.
The first one is Authority Information Access, or AIA. The AIA allows end users to obtain the certificate used by the CA itself. This can be very useful if you have a root CA that is offline for security reasons.
This is not an uncommon configuration. Some organizations will keep their root CAs offline all the time to make them less susceptible to compromise. But we still need to have a copy of the root CA certificate.
To make this possible, we use AIA to publish a copy of the root CA certificate to some other location so that users can still access the root CA certificate. This is important because we need that root CA certificate with the CA's public key in order to verify digital signatures that the CA has implemented.
If the CA digitally signs something, you must have a copy of its public key to verify that signature. AIA makes that public key available offline. If we didn't do this and the root CA was offline, we'd have to figure out some other way to get a copy of that key.
In addition to AIA, we also have CRL and CRL Distribution Point, also known as CDP. This is where users go to get a copy of the latest CRL. The same issue exists here as exists with the root CA's public key.
If the root CA is offline, we can't access the CRL. We can't let everybody know which certificates have been revoked. What we can do is publish the CRL to some other location that clients can access.
For example, we can publish it to Active Directory via LDAP. We could access it on a web page, or we could put it in the file system.
Let's go ahead and click Cancel. If we go back to our CertSRV web page, we can see there's an option to download the CA certificate or view the certificate chain or CRL.
This option, Download a CA certificate, is AIA. This one is CDP. If we click this option, we could download the latest CRL or the latest delta CRL.
In this demonstration, we covered various aspects of managing certificates within the public key infrastructure. We discussed requesting certificates, manual approval, issuing certificates, and revoking certificates.
We also delved into the availability of the CA's public keys via AIA and the publication of the Certificate Revocation List (CRL) to inform users about revoked certificates.

3.5.6 Manage Certificates (Simulation)

Scenario

You are the IT administrator for a growing corporate network. You manage the certification authority for your network. As part of your daily routine, you perform several certificate management tasks. CorpCA, the certification authority, is a guest server on CorpServer2.

In this lab, your task is to complete the following:

- Your network uses smart cards to control access to sensitive computers. Currently, the approval process dictates that you manually approve smart card certificate requests.
  Approve pending certificate requests for smart card certificates from tsutton and mmallory.
- Deny the pending web server certificate request for CorpSrv12.
- User bchan lost his smartcard. Revoke the certificate assigned to bchan.CorpNet.com using the **Key Compromise** reason code.
- Unrevoke the CorpDev3 certificate.

Explanation

Complete this lab as follows:

1. Access Certification Authority on the CORPSERVER2 server.
   a. From Hyper-V Manager, select **CORPSERVER2** .
   b. Maximize the window for better viewing.
   c. From the Virtual Machines pane, double-click **CorpCA** .
   d. From the Server Manager's menu bar, select **Tools > Certification Authority** .
   e. Maximize the window for better viewing.
   f. From the left pane, expand **CorpCA-CA** .
2. Approve the pending certificate request for tsutton and mmallory.
   a. Select **Pending Requests** .
   b. From the right pane, scroll until you can see the **Requester Name** column.
   c. Right-click on the row that contains **tsutton** and select **All Tasks > Issue** to approve the certificate.
   d. Right-click on the row that contains **mmallory** and select **All Tasks > Issue** .
3. Deny the pending request for CorpSrv12.
   a. Right-click on the row that contains **CorpSrv12.CorpNet.com** and select **All Tasks** > **Deny** .
   b. Select **Yes** to confirm the denial.
4. Revoke bchan's certificates.
   a. From the left pane, select **Issued Certificates** .
   b. From the right pane, right-click **bchan.CorpNet.com** and select **All Tasks** > **Revoke Certificate** .
   c. Using the *Reason code* drop-down menu list, select **Key Compromise** .
   d. Select **Yes** .
5. Unrevoke the CorpDev3 certificate.
   a. From the left pane, select **Revoked Certificates** .
   b. From the right pane, right-click **CorpDev3.CorpNet.com** and select **All Tasks** > **Unrevoke Certificate** .

3.5.7 Certificate Concepts (Lesson Video)

Transcript:

Using digital certificates to share public keys and validate organizations has become a key component to doing business over the internet. Certificate Authorities, or CAs, are trusted organizations that validate and administer these certificates. In this lesson, I'll go over how CAs are set up and managed.

To fully understand how certificates work, you need to understand how CAs are set up. To increase security, CAs operate in a hierarchy of multiple CAs. This is done so that if one of the CAs is compromised, only the certificates it issued need to be revalidated.
The first CA is the Root CA. This is a self-signed certificate and is used to validate additional CAs.
Under the Root CA, you have Subordinate CAs, also known as Intermediate CAs. You can have multiple Intermediate CAs based on their policies and regulations. The Intermediate CAs validate the Issuing Certificate Authorities and the

Issuing CA is the one that hands out certificates. When you check the certification path in a web browser, you can usually see this structure.
For example, when we look at TestOut's certificate, we see that their server certificate was issued by this Intermediate CA, which was signed by this Root CA. This structure is known as certificate chaining, or the chain of trust.
Certificate chaining is done to protect the root certificate. If the root certificate is ever compromised, all certificates that have been issued underneath it are now invalid and need to be redone.
Another common method of protecting the Root CA is to keep it offline and only bring it online when it needs to authorize a new Intermediate CA.
Once the Root CA has been set up and the intermediate certificates issued, the Root CA is generally taken offline. This means that it's isolated from network access and is usually completely turned off. If an Intermediate CA is compromised, only the certificates they issued need to be reissued. In other words, not every certificate that was issued under the Root CA is compromised.
The problem with this is that the Root CA can't manage and update the Certificate Revocation List, or CRL, if it's offline. To remedy this, the Intermediate CAs can be set up to manage the CRL, or a specific Intermediate CA can be configured to handle this job.
Instead of having browsers constantly checking the CRLs, the Online Certificate Status Protocol, or OCSP, was created to replace CRLs and can be used to determine the validity of certificates. Let's see how this works.
When OCSP servers, also called responders, receive a request from the web browser, they check with the CA and send an up-to-date certificate status. This frees up CAs from having to maintain their own CRLs.
To further help with performance, a network might use OCSP stapling. Stapling simply means that the server that holds the certificate also provides any revocation information.
This works by having the certificate holder send a query to the OCSP responder at set intervals. Then the server staples that timestamped response to its certificate so that in the first SSL or TSL handshake the OCSP validation is already attached. Now, the user's browser doesn't need to perform a separate OCSP request.
Keep in mind that it's possible for a man-in-the-middle attack to intercept certificates during the handshake. This would allow a hacker to set up a malicious site using a valid certificate. This can be prevented using certificate pinning.
Certificate pinning is when an application, such as a web browser, has the server's certificate hard coded in the browser itself. When the application connects to the server, it downloads and checks the certificates. If they don't match, the application blocks the connection. Just know that certificate pinning was never really adopted by modern browsers, so it would mostly be used inside of an organization that's running their own CA.
Let's shift gears to trust models. All CAs start with a single trust model. This means that there's one Root CA for the organization that issues and manages certificates. As you've already seen, a hierarchical structure is mainly used. It starts with the root certificate and multiple Intermediate CAs below it.
It's also possible for you to set up a mesh model. In a mesh model, multiple CAs issue certificates to each other. This works great because if one CA is compromised, those certificates can still be valid because other CAs have authenticated them. The problem with this model is that it becomes extremely difficult to scale up with growth.
Another thing to consider is how trusts can be set up when one organization who has their own CA structure wants to partner with another organization that has their own CA structure. In this scenario, there needs to be a way for each organization to trust the certificates issued by the other's CAs. This can be done by setting up a bridge model in which both hierarchies trust each other. In this case, clients in both organizations trust certificates issued by CAs of either organization. You can also be more restrictive by setting up trusts with individual CAs that are lower on the hierarchy.
Another common setup is the web of trust. Instead of using a CA, everyone is a trusted authority.
For example, if Craig trusts Ethan and Lynette trusts Craig, Lynette will also have an indirect trust with Ethan. In other words, you trust a friend of a friend. We see this setup often in PGP.
Implementing the proper model can help ensure that certificates are properly maintained and validated. Of course, all of this also relies on keeping our encryption keys safe.
Keeping a backup of private keys is important, but you'll need to make sure that they can't be accessed unless absolutely necessary. Normally, the CA never gets the user's private keys.
You can, though, use key archival to back up private keys. All this means is that the user sends their private key in a secure transmission to the CA to keep safe. The problem with this is that a serious security breach can occur if private keys are compromised. The advantage is that the archive of private keys is readily available if any are lost or corrupted.
Key escrow is like key archival, except that keys are sent to a trusted third party instead of the CA holding the keys. This is often done for security or legal purposes. Getting access to these keys may require legal action, and this is by design. For example, if law enforcement needs to access and analyze emails, they would need a court order to get the keys needed.
That's it for this lesson. In this lesson, we discussed how CAs are set up to manage certificates. We first looked at the concept of certificate chaining, which is the method of using a hierarchical certificate structure. Then we covered the

difference in online and offline CAs and the best practice of leaving them offline whenever possible. Finally, we discussed different trust models and how to assure key safety through key archival or escrow services.

3.5.8 Certificate Concepts Facts

Using digital certificates to share public keys and validate organizations is a critical component of doing business over the internet. Certificate authorities (CAs) are the trusted organizations that validate and administer digital certificates.

This lesson covers the following topics:

- Certificate revocation
- Root of Trust
- Single CA
- Third-party CAs
- Private key safety

Certificate Revocation

A certificate may be revoked or suspended:

- A revoked certificate is no longer valid and cannot be "un-revoked" or reinstated.
- A suspended certificate can be re-enabled.

A certificate may be revoked or suspended by the owner or the CA for many reasons. For example, the private key may have been compromised, or the business could have closed, a user could have left the company, a domain name could have been changed, the certificate could have been misused, and so on. These reasons are codified under choices such as Unspecified, Key Compromise, CA Compromise, Superseded, or Cessation of Operation. A suspended key is given the code Certificate Hold.

There must be a mechanism to inform users whether a certificate is valid, revoked, or suspended. A CA must maintain a **certificate revocation list (CRL)** of all revoked and suspended certificates. The CRL must be accessible to anyone relying on the validity of the CA's certificates. Each certificate should contain information for the browser on how to check the CRL.

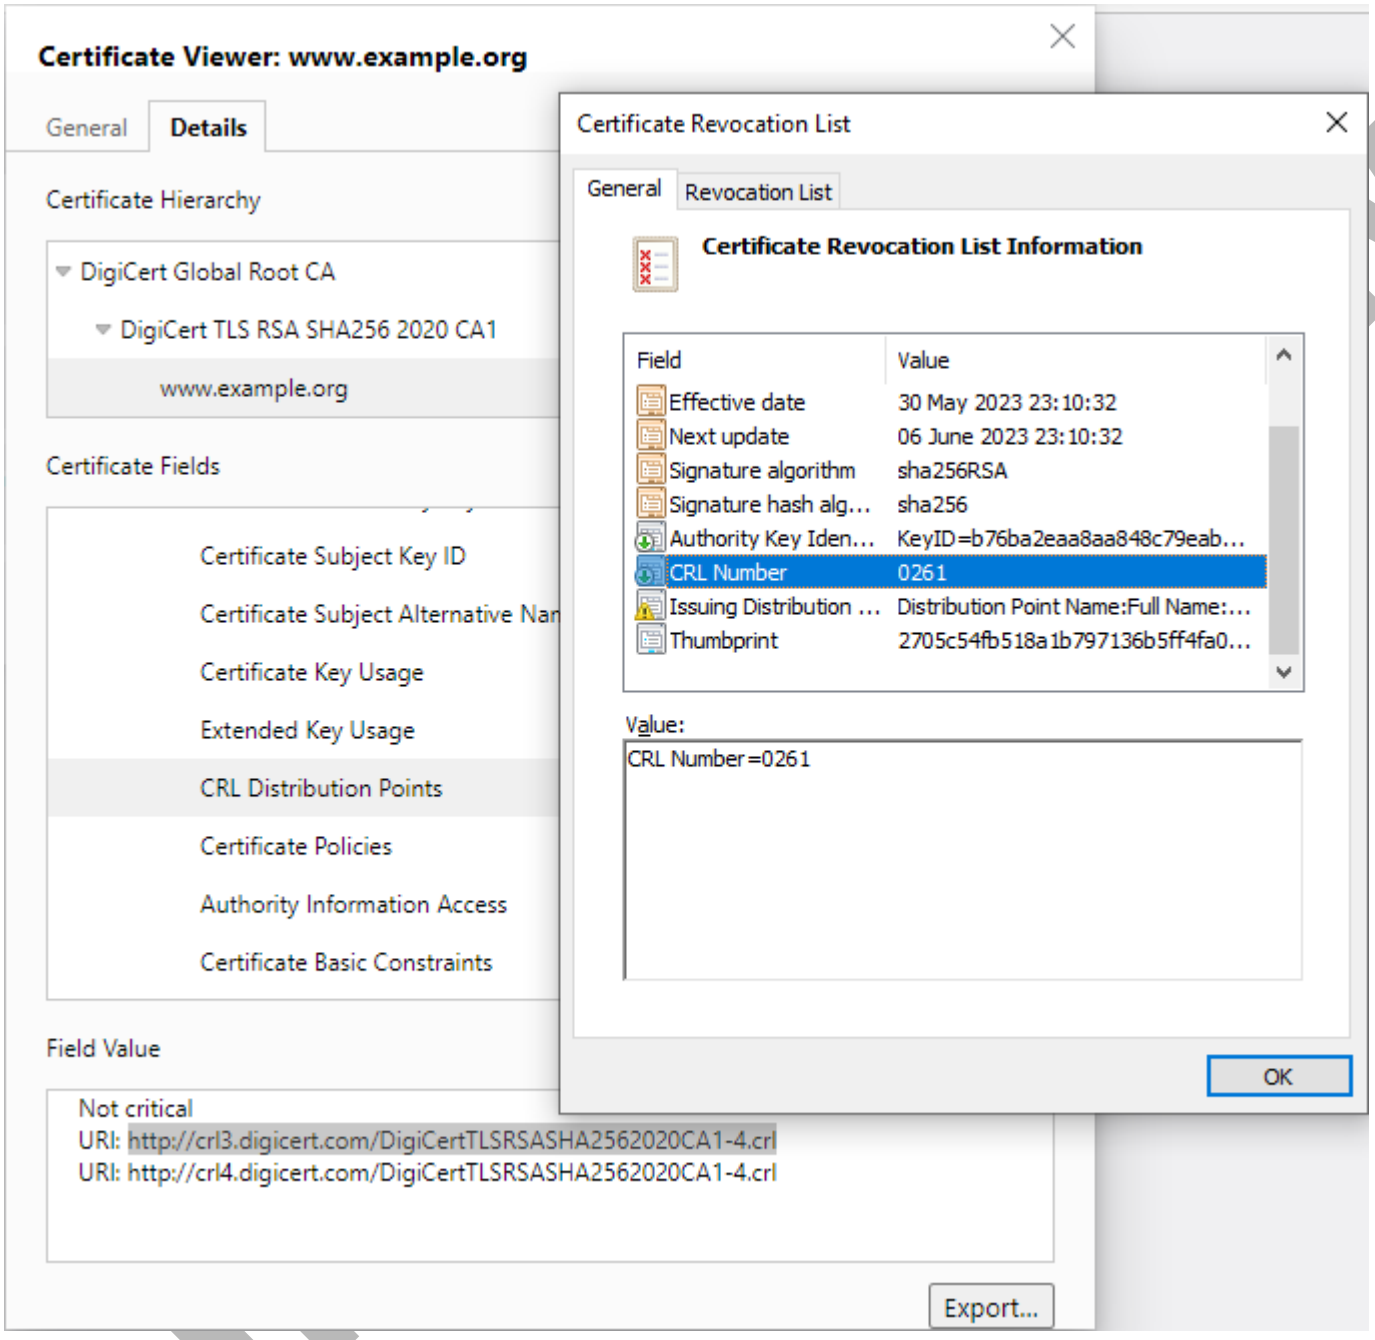

The distribution point field in a digital certificate identifies the location of the list of revoked certificates, which are published in a CRL file signed by the CA. (Screenshot used with permission from Microsoft.)

A CRL has the following attributes:

- **Publish Period** - the date and time on which the CRL is published. Most CAs are set up to publish the CRL automatically.
- **Distribution Point(s)** - the location(s) to which the CRL is published.
- **Validity Period** - the period during which the CRL is considered authoritative. This is usually a bit longer than the publish period (for example, if the publish period was every 24 hours, the validity period might be 25 hours).

- **Signature** - the CRL is signed by the CA to verify its authenticity.

With the CRL system, there is a risk that the certificate might be revoked but still accepted by clients because an up-to-date CRL has not been published. A further problem is that the browser (or other application) may not be configured to perform CRL checking, although this now tends to be the case only with legacy browser software.

Another means of providing up-to-date information is to check the certificate's status on an Online Certificate Status Protocol (OCSP) server. Rather than return a whole CRL, this communicates the requested certificate's status. Details of the OCSP responder service should be published in the certificate.

Most OCSP servers can query the certificate database directly and obtain the real-time status of a certificate. Other OCSP servers depend on the CRLs and are limited by the CRL publishing interval.

Root of Trust

The *root of trust model* defines how users and different CAs can trust one another. Each CA issues itself a certificate. This is referred to as the root certificate . The root certificate is self-signed, meaning the CA server signs a certificate issued to itself. A root certificate uses an RSA key size of 2,048 or 4,096 bits or the ECC equivalent. The subject of the root certificate is set to the organization/CA name, such as "CompTIA Root CA."

The root certificate can be used to sign other certificates issued by the CA. Installing the CA's root certificate means that hosts will automatically trust any certificates signed by that CA.

Single CA

In this simple model, a single root CA issues certificates directly to users and computers. This single CA model is often used on private networks. The problem with this approach is that the single CA server is very exposed. If it is compromised, the whole PKI collapses.

Third-party CAs

Most third-party CAs operate a hierarchical model. In the hierarchical model, the root CA issues certificates to one or more intermediate CAs. The intermediate CAs issue certificates to subjects (leaf or end entities). This model has the advantage that different intermediate CAs can be set up with certificate policies, enabling users to perceive clearly what a particular certificate is designed for. Each leaf certificate can be traced to the root CA along the certification path. This is also referred to as certificate chaining or a *chain of trust* .

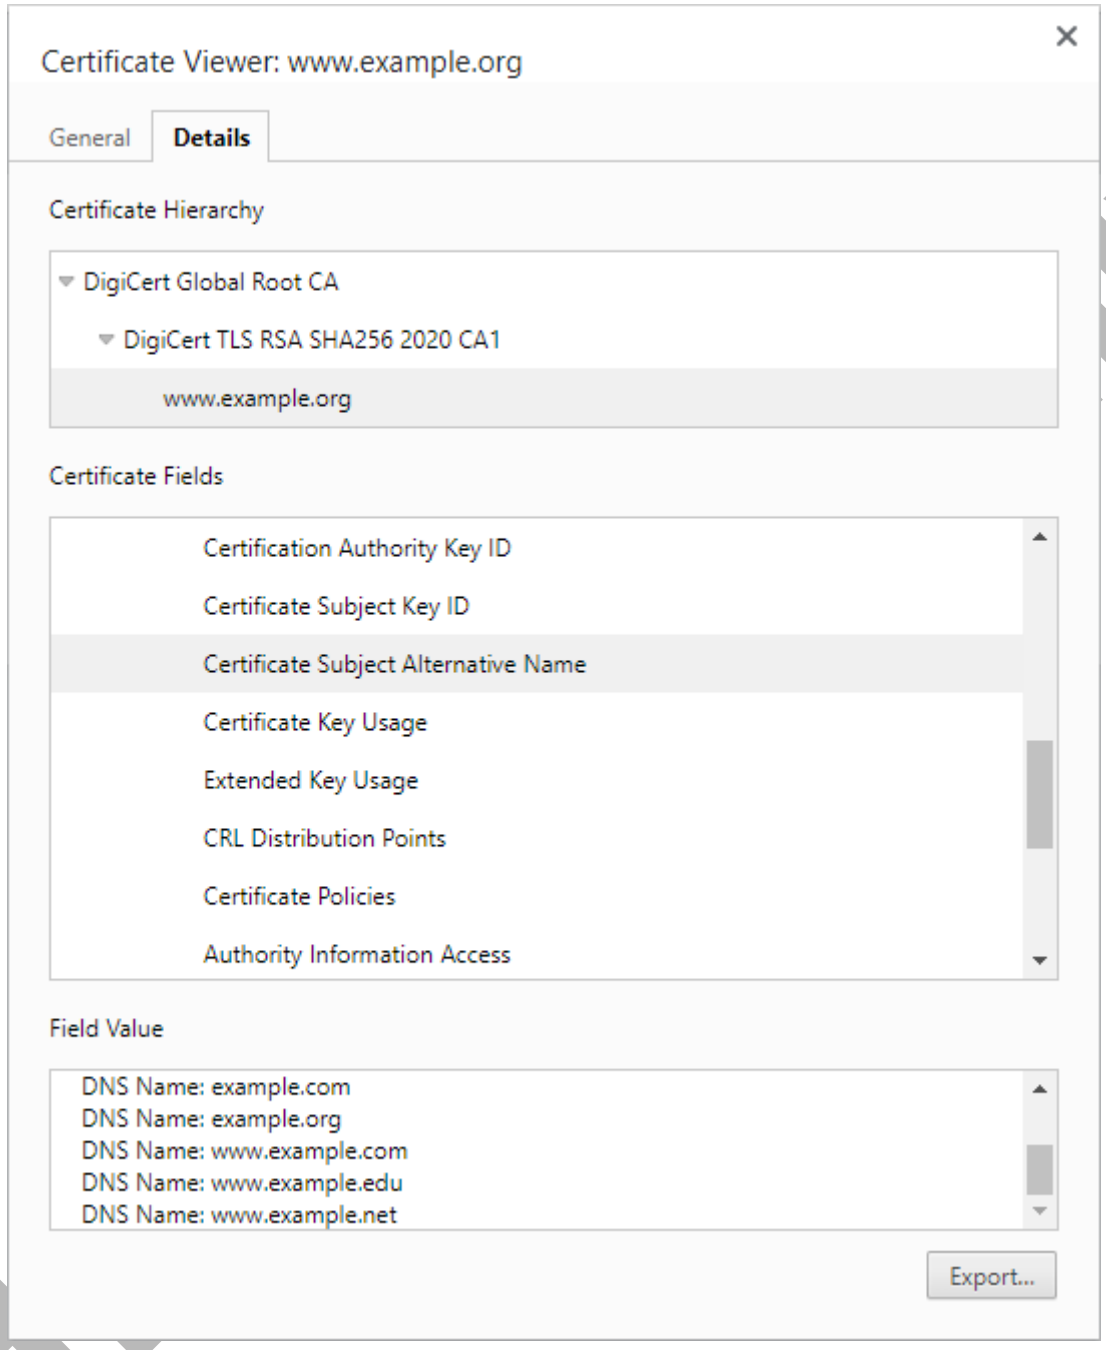

The web server for www.example.org is identified by a certificate issued by the DigiCert TLS CA1 intermediate CA. The intermediate CA's certificate is signed by DigiCert's Global Root CA (Screenshot used with permission from Microsoft).

In some circumstances, using PKI can be too difficult or expensive to manage. Any machine, web server, or program code can be deployed with a self-signed certificate . For example, the web administrative interfaces of consumer routers are often only protected by a self-signed certificate. Self-signed certificates can also be useful in development and test environments. The operating system or browser will mark self-signed certificates as untrusted, but a user can choose to override this. The nature of self-signed certificates makes them very difficult to validate. They should not be used to protect critical hosts and applications.

Private Key Safety

To ensure data can always be recovered, you should create a backup of the private keys. It is important to have a backup and, equally important, that the backup is kept safe. The following table shows two main methods to backup private keys:

Key Backup Method	Description
Key archival	In key archival, the key is backed up by the CA. To do this, the user sends the private key in a secure transmission to the CA to back it up. This method is often used in an organization that manages its own CA. If keys are lost, they will be readily available and easily accessed. However, if the CA is breached, all private keys will be compromised.
Key escrow	If a private or secret key is lost or damaged, ciphertexts cannot be recovered unless a backup of the key has been made. Making copies of the key is problematic as it becomes more likely that a copy will be compromised and more difficult to detect that a compromise has occurred. These issues can be mitigated by using escrow and M of N controls. Escrow means that something is held independently. In terms of key management, this refers to archiving a key (or keys) with a third party. M of N means that an operation cannot be performed by a single individual. Instead, a quorum (M) of available persons (N) must agree to authorize the operation. A key can be split into one or more parts. Each part can be held by separate escrow providers, reducing the risk of compromise. An account with permission to access a key held in escrow is referred to as a key recovery agent (KRA). A recovery policy can require two or more KRAs to authorize the operation. This mitigates the risk of a KRA attempting to impersonate the key owner.

3.5.9 Certificates and Certificate Authorities

3.5.10 Practice Questions (Section Quiz)

q_cryt_pki_ca_01_secp8

Which aspect of a certificate makes it a reliable and useful mechanism for proving the identity of a person, system, or service on the internet?

Answers:

- ***It is a trusted third party.**
- It uses electronic signatures.
- It provides ease of use.
- It is a digital mechanism rather than a physical one.

Explanation:

The use of a trusted third party (called a certificate authority or CA) is what makes certificates a reliable and useful mechanism for proving the identity of a person, system, or service on the internet. The CA issues proof of identity to each organization in the form of a certificate. The fact that all entities trust the CA makes the certificates trusted and valuable.

A certificate only proves identity; it does not prove reliability. Electronic signatures are a form of certificate that verifies identity. While electronic signatures prove identity, they do so only because both parties trust the authority of the CA, not because the

signature exists. It is true that certificates are easy to use. However, ease of use does not make them reliable. Certificates are a digital mechanism, which makes them suited for use on the internet. However, that alone does not make them reliable or useful.

q_cryt_pki_ca_02_secp8

An SSL client has determined that the certificate authority (CA) issuing a server's certificate is on its list of trusted CAs.

What is the next step in verifying the server's identity?

Answers:

- The domain on the server certificate must match the CA's domain name.
- The post-master secret must initiate subsequent communication.
- ***The CA's public key validates the CA's digital signature on the server certificate.**
- The master secret is generated from common key code.

Explanation:

Once an SSL client has identified a CA as trusted, it uses the CA's public key to validate the CA's digital signature on the server certificate. If the digital signature can be verified, the client accepts the server certificate as a valid certificate issued by a trusted CA.

SSL clients verify a server's identity using the following steps:

1. The client checks the server's certificate validity period. The authentication process stops if the current date and time fall outside of the validity period.
2. The client verifies that the issuing certificate authority is on its list of trusted CAs.
3. The client uses the CA's public key to validate the CA's digital signature on the server certificate. If the digital signature can be verified, the client accepts the server certificate as a valid certificate issued by a trusted CA.
4. To protect against man-in-the-middle attacks, the client compares the actual DNS name of the server to the DNS name on the certificate.

q_cryt_pki_crl_01_secp8

Certificates can be invalidated by the trusted third party that originally issued the certificate.

What is the name of the mechanism that is used to distribute information about invalid certificates?

Answers:

- ACL
- TACACS
- ***CRL**
- One-way function

Explanation:

The CRL (certificate revocation list) is the mechanism that is used to distribute information about invalid certificates. Each time an application receives a certificate, that application checks the CRL from the certificate authority (CA) that issued the certificate. If the certificate is not on the CRL and its timestamp is still valid, the user is prompted whether or not to accept the certificate.

ACLs are used to protect files and other resources.

TACACS is a remote access centralized authentication system.

One-way functions are common cryptographic mechanisms. None of these technologies are directly used to distribute invalid certificate information.

q_cryt_pki_crl_02_secp8

Which of the following BEST describes the contents of the CRL?

Answers:

- The current status of all certificates issued by a CA
- A list of all expired and revoked certificates
- ***A list of all revoked certificates**
- The archived private keys of all issued certificates

Explanation:

The certificate revocation list (CRL) resides at the CA and consists of a list of certificates that have been previously revoked.

Expired certificates are not put on the CRL; they are automatically invalid because the certificate validity period has passed. Archived private keys are held in key escrow, which is a third party that is trusted to protect the private keys.

q_cryt_pki_crl_03_secp8

Which of the following would require that a certificate be placed on the CRL?

Answers:

- The certificate validity period is exceeded.
- ***The private key is compromised.**
- The signature key size is revealed.
- The encryption key algorithm is revealed.

Explanation:

Certificates are published to the certificate revocation list (CRL) when a condition compromises the integrity of the certificate. If the private key is compromised (discovered), the certificate is no longer proof of identity.

Certificates do not need to be placed on the CRL if their validity period expires. In this case, the certificate simply expires. Knowing the signature key size or the encryption key algorithm does not compromise the integrity of the certificate.

q_cryt_pki_digital_certificates_secp8

You are a security analyst at a large corporation. The company uses digital certificates for secure communication.

One day, you notice that a certificate from a trusted third-party certificate authority (CA) has been flagged as invalid. The certificate is not expired, and the organization it was issued to still exists.

What could be the reason for this, and what should be your next step?

Answers:

- ***The private key associated with the certificate has been compromised. You should immediately inform your supervisor and initiate the process to obtain a new certificate.**
- The certificate was mistakenly flagged as invalid. You should ignore the warning and continue using the certificate.
- The certificate is a fake. You should report the issue to the CA and request a new certificate.
- The certificate was issued by a CA that is no longer trusted. You should replace the certificate with one issued by a trusted CA.

Explanation:

Immediately informing your supervisor and initiating the process to obtain a new certificate is the correct answer. If a certificate is flagged as invalid and the organization it was issued to still exists, it's possible that the private key associated with the certificate has been compromised. In this case, the certificate should be replaced immediately to prevent any potential security breaches.

Ignoring a warning about an invalid certificate could lead to serious security issues, such as data breaches or loss of sensitive information. It's important to investigate and resolve any issues related to digital certificates.

If the certificate was a fake, it would likely not have been accepted by the system in the first place, as it would fail the validation process. However, if this is the case, the issue should be reported to the CA and a new certificate should be requested.

If the CA that issued the certificate is no longer trusted, the certificate would not be flagged as invalid; instead, the system would not accept the certificate at all. If this were the case, the certificate would need to be replaced with one issued by a trusted CA.

q_cryt_pki_ocsp_secp8

Which technology was developed to help improve the efficiency and reliability of checking the validity status of certificates in large, complex environments?

Answers:

- Certificate revocation list
- Key escrow
- Private key recovery
- ***Online Certificate Status Protocol**

Explanation:

Online Certificate Status Protocol (OCSP) is the technology developed to improve the efficiency and reliability of checking the validity status of certificates in large, complex environments. OCSP allows clients to query a CA or registration authority (RA) and quickly learn whether a certificate is valid or has been revoked.

OCSP is a significant improvement over the CRL mechanism. CRLs were static lists that were distributed periodically to CAs and RAs. However, CRLs were often out of date. Key escrow and private key recovery are not related to certificate status checking.

q_cryt_pki_pki_01_secp8

Which of the following is a mechanism for granting and validating certificates?

Answers:

- ***PKI**
- RADIUS
- Kerberos
- AAA

Explanation:

Certificates are obtained from a public key infrastructure (PKI). A PKI is a system that allows a trusted third party to vouch for user identities. A PKI is made up of certificate authorities (CAs), while a CA is an entity trusted to issue, store, and revoke certificates.

Both RADIUS and TACACS+ are protocols used for centralized authentication, authorization, and accounting with remote access. Kerberos is an authentication and authorization program that uses tickets.

q_cryt_pki_pki_02_secp8

A PKI is an implementation for managing which type of encryption?

Answers:

- ***Asymmetric**
- Symmetric
- Hashing
- Steganography

Explanation:

A public key infrastructure (PKI) is a hierarchy of computers that issue and manage certificates. Certificates use asymmetric encryption with a public and private key pair.

Because certificates use asymmetric encryption, and certificates are a key component of PKI, symmetric encryption is not utilized with PKI.

Hashing is the practice of transforming a given key or string of characters into another value for the purpose of security. Hashing is always used for one-way encryption, which makes it a symmetric (not asymmetric) encryption type.

Steganography is a type of encryption where digital content is hidden within another message or physical object to avoid detection. Steganography can be used to hide virtually any type of digital content, including text, image, video, or audio content. This is a type of symmetric encryption that is not part of the PKI process.

q_cryt_pki_purpose_secp8

An organization has tasked a cyber security technician with enhancing its framework after recently experiencing a cyber breach.

What is the value associated with a public key infrastructure (PKI)?

Answers:

- ***It is the framework that establishes trust in using public key cryptography to sign and encrypt messages via digital signatures.**
- It is a cryptoprocessor implemented as a module within the CPU on a computer or mobile device.
- It is a cryptoprocessor that implements hardware through a removable or dedicated form factor, such as plug-in PCIe adaptor cards.
- It summates all revoked and suspended certificates and must be accessible to anyone relying on the validity of the CA's certificates.

Explanation:

In this scenario, a PKI is the preferred framework that helps establish trust in using public key cryptography to sign and encrypt messages via digital signatures.

Not viable in this scenario, a trusted platform module (TPM) is a cryptoprocessor implemented as a module within the CPU on a computer or mobile device.

A hardware security module (HSM) is a cryptoproccessor that implements hardware through a removable or dedicated form factor, such as plug-in peripheral component interconnect express (PCIe) adaptor cards.

A certificate revocation list (CRL) provides a summation of all revoked and suspended certificates and must be accessible to anyone relying on the validity of the certificate authority's (CA's) certificates.

q_cryt_pki_ra_01_secp8

Which of the following is an entity that accepts and validates information contained within a request for a certificate?

Answers:

- ***Registration authority**
- Certificate authority
- Recovery agent
- Enrollment agent

Explanation:

A registration authority (RA) can be used in large enterprise environments to offload client enrollment request processing by handling client verification prior to certificates being issued. The RA accepts registrations, validates identity, and approves or denies certificate requests.

A certificate authority (CA) is an entity trusted to issue, store, and revoke digital certificates. Often, the role of a CA is combined with that of RA. But technically speaking, a CA is the computer that issues the certificate.

Recovery agents are users who are given the ability to restore private keys from the archive.

An enrollment agent is someone who can request a certificate on behalf of another user. Enrollment agents are often used to request certificates used on smart cards.

q_cryt_pki_ra_02_secp8

In the process of obtaining a digital certificate, which entity may a certificate authority rely on to perform the validation of the certificate signing request (CSR)?

Answers:

- Root authority
- ***Registration authority**
- Certificate revocation list
- Online Certificate Status Protocol

Explanation:

Registration authority is the correct answer. Registration authority (RA) is an entity that is certified by a root certificate authority and is authorized to issue certificates for specific uses only. The RA can perform the validation of the certificate signing request (CSR) on behalf of the certificate authority.

A root authority, or root CA, is the top-level certificate authority in a certificate chain. It is responsible for issuing certificates to subordinate CAs, not for validating CSRs.

A Certificate revocation list (CRL) is a list of certificates that have been revoked before their scheduled expiration date and are no longer valid. It does not perform any validation of CSRs.

The Online Certificate Status Protocol (OCSP) is an internet protocol used to obtain the revocation status of a digital certificate. It does not perform any validation of CSRs.

q_cert_types_ca_secp8

In the certificate authority trust model known as a hierarchy, where does trust start?

Answers:

- ***Root CA**
- Issuing CA
- Third-party CA
- Registration authority

Explanation:

Trust starts at the Root CA in all trust models.

An Issuing CA can be a Root CA or a CA at any level below the root.

A third-party CA may be the source of trust, but even then, the trust starts at a Root CA located somewhere.

A registration authority (RA) is a limited-functionality CA where certificates are verified, but no new certificates can be issued.

q_cert_types_code_cert_secp8

Sarah is a software developer for a tech startup, TechInnovate. She has developed a new application that the company plans to distribute to its customers.

Sarah wants to ensure that the application is seen as legitimate and trustworthy by the users and their systems. She also wants to protect the integrity of the application code from being tampered with.

Which type of certificate should Sarah use?

Answers:

- Root certificate
- Subject Alternative Name (SAN) certificate
- ***Code signing certificate**
- Self-signed certificate

Explanation:

A code signing certificate is used by software developers to digitally sign their applications. This proves the legitimacy of the application and ensures that the code has not been tampered with since it was signed. This is exactly what Sarah needs for her application.

A root certificate is the top-most certificate in a certificate chain, used by certificate authorities (CAs) to sign other certificates. It doesn't serve the purpose of validating software applications or protecting their integrity.

A Subject Alternative Name (SAN) certificate is used to secure multiple domains with a single certificate. It doesn't provide the functionality to validate software applications or protect their integrity.

A self-signed certificate is not issued by a trusted CA and would likely cause trust issues with systems and users. It also doesn't provide the functionality to validate software applications or protect their integrity.

q_cert_types_csr_01_secp8

You are a security analyst at a large organization. Your organization uses a third-party certificate authority (CA) for its public key infrastructure (PKI).

One day, you receive a certificate signing request (CSR) from a new department in your organization. The CSR contains a public key and the department's information. However, you notice that the department's information does not match the information in the organization's official records.

What should you do?

Answers:

- Approve the CSR and issue the certificate.
- Reject the CSR and ask the department to submit a new one with correct information.
- Ignore the discrepancy and forward the CSR to the third-party CA.
- ***Investigate the discrepancy and verify the department's information before forwarding the CSR to the third-party CA.**

Explanation:

Investigating the discrepancy and verifying the department's information before forwarding the CSR to the third-party CA is the correct answer. As a security analyst, it's your responsibility to ensure the integrity of the organization's PKI. This includes verifying the information in CSRs before they are forwarded to the CA. Investigating the discrepancy and verifying the department's information ensures the CSR is legitimate and the certificate issued will be valid.

Approving the CSR and issuing the certificate without verifying the department's information can lead to security risks. The certificate could be used for malicious purposes.

While it's important to ensure the CSR contains correct information, simply rejecting the CSR without investigating the discrepancy could overlook potential security issues.

Forwarding the CSR to the third-party CA without addressing the discrepancy could result in the issuance of a certificate with incorrect information. This could lead to security vulnerabilities.

q_cert_types_csr_02_secp8

You are a network administrator for a large corporation. You receive a Certificate Signing Request (CSR) from a department within your organization. The CSR is for a new server that will be used to host sensitive company data.

The department has requested a Subject Alternative Name (SAN) certificate to cover multiple domains. However, you notice that one of the domains listed in the CSR is not owned by your organization.

What should you do?

Answers:

- Issue the SAN certificate as requested.
- ***Reject the CSR and ask the department to submit a new one without the unowned domain.**
- Issue the SAN certificate but exclude the unowned domain.
- Ignore the unowned domain and forward the CSR to the certificate authority (CA).

Explanation:

The CSR should be rejected and the department should be asked to submit a new one without the unowned domain. This ensures that the issued certificate only covers domains owned by your organization.

Issuing the SAN certificate as requested would mean including a domain not owned by your organization. This could lead to security risks and legal issues.

While it might seem like a good idea to issue the certificate and exclude the unowned domain, this could lead to confusion and potential security risks in the future. It's better to reject the CSR and ask for a new one to ensure the request is accurate and complete.

Ignoring the unowned domain and forwarding the CSR to the CA could result in the issuance of a certificate that includes a domain not owned by your organization. This could lead to security vulnerabilities and potential legal issues.

q_cert_types_org_validation_secp8

A medium-sized e-commerce company is planning to upgrade their website's security by acquiring a certificate from a certificate authority (CA).

The company wants to ensure that the certificate not only validates their domain ownership but also verifies the legitimacy of their organization. They are also looking for a validation process that can be completed within 1 to 3 days.

As the IT manager for the company, which level of CA validation would you recommend?

Answers:

- Domain validation
- ***Organization validation**
- Extended validation
- Self-signed certificate

Explanation:

Organization validation is the best choice because it requires the purchaser to prove they are a domain administrator and also prove the organization is legitimate. Additionally, these certificates can be issued in 1-3 days, which aligns with the company's requirement.

Domain validation only validates domain ownership and does not verify the legitimacy of the organization.

Extended validation provides a thorough validation process but it can take up to 5 days to be issued, which is longer than the company's preferred timeline.

Self-signed certificate does not provide the same protection and security as a CA-validated certificate and is not recommended for an e-commerce company like the one in the scenario.

q_cert_types_san_cert_secp8

John is a network administrator for a growing company, XYZ Corp, which has recently acquired two smaller companies, ABC Corp and DEF Corp.

Each company has its own domain: xyz.com, abc.net, and def.org. John needs to secure all three domains with a single certificate to simplify management and reduce costs. He also wants to ensure that the certificate is issued by a trusted certificate authority (CA) to avoid trust issues with web browsers.

Which type of certificate should John choose?

Answers:

- Root certificate
- ***Subject Alternative Name (SAN) certificate**
- Wildcard certificate
- Self-signed certificate

Explanation:

A Subject Alternative Name (SAN) certificate is the correct answer. A SAN certificate allows an organization to cover multiple domains with one certificate. This is exactly what John needs to secure all three domains (xyz.com, abc.net, def.org) with a single certificate.

A root certificate is the first certificate that a certificate authority (CA) creates and is used to sign lower-level certificates. It is not suitable for John's needs as it does not cover multiple domains.

While a wildcard certificate can cover one domain and multiple subdomains, it cannot cover multiple domains. Therefore, it would not meet John's needs to secure all three domains with one certificate.

A self-signed certificate is not issued by a trusted CA and would likely cause trust issues with web browsers. This would not meet John's requirement for the certificate to be issued by a trusted CA.

q_cert_types_wildcard_secp8

The network administrator for an international e-commerce company that operates multiple online stores must ensure secure communication across various subdomains.

To streamline secure sockets layer/transport layer security (SSL/TLS) certificate management and implement a robust public key infrastructure (PKI), the network administrator must identify the most suitable solution for efficiently securing the company's numerous subdomains within the PKI.

What is the MOST suitable solution for efficiently securing the multiple subdomains of the company's online stores within the PKI?

Answers:

- ***Wildcard certificates**
- Self-signed certificates
- Certificate revocation lists (CRLs)
- Certificate pinning

Explanation:

By using wildcard certificates, the company can secure all subdomains under a single certificate, ensuring efficient management and reducing administrative overhead. Wildcard certificates streamline the certificate deployment process and simplify ongoing maintenance tasks, leading to improved operational efficiency.

Self-signed certificates lack the trust and validation provided by a reputable certificate authority.

Certificate revocation lists (CRLs) provide information about revoked certificates but do not address the efficient management of multiple subdomains.

Certificate pinning is a mechanism to ensure trust in specific certificates but not specifically designed to manage multiple subdomains within a PKI efficiently.

q_cert_concepts_crl_01_secp8

A private key has been stolen. Which action should you take to deal with this crisis?

Answers:

- Delete the public key
- ***Add the digital certificate to the CRL**
- Place the private key in escrow
- Recover the private key from escrow

Explanation:

If a private key--a digital certificate or digital signature--is compromised (especially by theft), it should be added to the CRL. This prevents any future use of the key/certificate and prevents impersonation attacks.

There is no need to delete the public key because CRLs deal with any attempted use of the private key. The private key should have been placed in escrow at the beginning of its lifetime if key recovery was desired. In this situation, key recovery is not necessary.

q_cert_concepts_crl_02_secp8

Which action is taken when the private key associated with a digital certificate becomes compromised?

Answers:

- ***The compromised digital certificate is added to the certificate revocation list (CRL).**
- The compromised digital certificate is reissued with the same private key.
- The compromised digital certificate is ignored and continues to be used.
- The compromised digital certificate is deleted from the system.

Explanation:

The correct answer is that the compromised digital certificate is added to the certificate revocation list (CRL). A CRL is a list of digital certificates that have been revoked by the issuing certificate authority (CA) before their scheduled expiration date and should no longer be trusted. When a private key associated with a digital certificate becomes compromised, the certificate is added to the CRL to ensure that it is no longer trusted by any systems or users.

Reissuing the compromised digital certificate with the same private key would not solve the problem. The private key has been compromised, meaning unauthorized individuals may have access to it. Using the same private key would continue to leave the certificate and any associated systems or data vulnerable to misuse.

Ignoring the compromise and continuing to use the digital certificate would be a serious security risk. The compromised private key could be used by unauthorized individuals to impersonate the certificate holder, decrypt sensitive data, or perform other malicious activities.

Simply deleting the compromised digital certificate from the system would not be sufficient to address the issue. Other systems or users that trust the certificate would not be aware of the compromise and could continue to trust the certificate if they encounter it. The certificate needs to be revoked and added to the CRL to ensure it is no longer trusted.

q_cert_concepts_crl_attributes_secp8

Which of the following are attributes of a certificate revocation list (CRL)? (Select two.)

Answers:

- ***Publish period**
- ***Distribution point(s)**
- Validity period
- Certificate authority
- Encryption algorithm

Explanation:

The following are attributes of a certificate revocation list (CRL):

- Publish period - This is the date and time on which the CRL is published. Most CAs are set up to publish the CRL automatically.
- Distribution point(s) - These are the location(s) to which the CRL is published.

The following are *not* attributes of a Certificate Revocation List (CRL):

- Certificate authority - While the CA is responsible for creating and managing the CRL, it is not an attribute of the CRL itself.
- Encryption algorithm - This is not an attribute of the CRL. It is a method used to secure data, not a characteristic of the CRL.

- Validity period - This is a common misconception. While it seems logical that a CRL would have a validity period, in reality, the CRL is continuously updated and does not have a set validity period. Instead, each individual certificate within the CRL has its own validity period.

q_cert_concepts_escrow_01_secp8

What is the purpose of key escrow?

Answers:

- ***Key escrow is a method of storing a copy of the encryption key with a trusted third party.**
- Key escrow is a method of generating encryption keys.
- Key escrow is a method of transmitting encryption keys over the internet.
- Key escrow is a method of revoking encryption keys.

Explanation:

The correct answer is that key escrow is a method of storing a copy of the encryption key with a trusted third party. Key escrow is a security measure where a copy of the encryption key is stored with a trusted third party. This is often done for security and legal purposes. In case the original key is lost or the key holder is unavailable, the key can be retrieved from the escrow to access the encrypted data.

Key escrow does not involve the generation of encryption keys. It is a method for storing a copy of the encryption key with a trusted third party for safekeeping and retrieval in case of need.

Key escrow is not a method for transmitting encryption keys. It is a method of storing a copy of the encryption key with a trusted third party. The transmission of encryption keys involves different protocols and methods to ensure security.

Key escrow does not involve the revocation of encryption keys. It is a method of storing a copy of the encryption key with a trusted third party. The revocation of encryption keys is a separate process that involves invalidating the key and ceasing its use for encryption or decryption.

q_cert_concepts_escrow_02_secp8

You are concerned that if a private key is lost, all documents encrypted with your private key will be inaccessible.

Which service should you use to solve this problem?

Answers:

- ***Key escrow**
- OCSP
- RA
- CSP

Explanation:

Key escrow backs up private keys to a third-party organization outside of the company. If the private key is lost, you can recover the key from escrow.

Online Certificate Status Protocol (OCSP) is a protocol used to check the status of an individual digital certificate to verify whether it is good or has been revoked.

A cryptographic service provider (CSP) resides on the client and generates the key pair.

A registration authority (RA) verifies the information included in a certificate request.

q_cert_concepts_escrow_03_secp8

A network administrator responsible for managing the encryption keys used in the organization's secure communications had a new key management policy implemented by the organization, which included a provision for key escrow.

The administrator understands the role of key escrows in relation to private keys.

Which of the following BEST describes the purpose of key escrow in the context of private keys?

Answers:

- ***Key escrow involves securely storing a copy of the private key with a trusted third party for recovery purposes, ensuring availability in case of key loss or compromise.**
- Key escrow refers to the process of encrypting the private key to prevent unauthorized access and protect it from disclosure.
- Key escrow enables the distribution of private keys to multiple users within the organization to enhance key redundancy and reliability.
- Key escrow involves periodically rotating the private key to maintain data integrity and prevent cryptographic attacks.

Explanation:

Key escrow involves securely storing a copy of the private key with a trusted third party for recovery and availability in case of key loss or compromise. Key escrow serves as a backup mechanism for private keys.

Key escrow does not primarily involve encrypting the private key. The primary purpose of key escrow is key recovery and availability.

Key escrow does not facilitate the distribution of private keys to multiple users for redundancy and reliability.

Key escrow does not involve the periodic rotation of private keys. The primary purpose does not relate to data integrity or preventing cryptographic attacks.

q_cert_concepts_root_of_trust_secp8

Which of the following statements accurately describes the root of trust model in a public key infrastructure (PKI)?

Answers:

- The root of trust model involves multiple root certificates, each issued by a different certificate authority (CA).
- In the root of trust model, the root certificate is issued by a third-party CA, not the organization's own CA.
- ***The root of trust model defines how users and different CAs can trust one another, with each CA issuing itself a root certificate.**
- The root of trust model involves a root certificate that is issued by a user, not a CA.

Explanation:

The root of trust model defines how users and different CAs can trust one another, with each CA issuing itself a root certificate is the correct answer. The root of trust model defines how users and different CAs can trust one another, with each CA issuing itself a root certificate. This is the core concept of the root of trust model, where the root certificate is self-signed by the CA, and installing the CA's root certificate means that hosts will automatically trust any certificates signed by that CA.

The root of trust model involves multiple root certificates, each issued by a different certificate authority (CA) is incorrect because in the root of trust model, there is typically one root certificate that is self-signed by the CA.

In the root of trust model, the root certificate is issued by a third-party CA, not the organization's own CA is incorrect because in the root of trust model, the root certificate is self-signed by the CA itself, not a third-party CA.

The root of trust model involves a root certificate that is issued by a user, not a CA is incorrect because in the root of trust model, the root certificate is issued by the CA, not a user.

q_cert_concepts_single_ca_secp8

Which of the following statements accurately describes the single certificate authority (CA) model in a public key infrastructure (PKI)?

Answers:

- In the single CA model, the root CA issues certificates to multiple intermediate CAs.
- The single CA model is often used in large, complex networks due to its scalability.
- ***In the single CA model, a single root CA issues certificates directly to users and computers.**
- The single CA model is highly secure and if the CA server is compromised, it does not affect the entire PKI.

Explanation:

In the single CA model, a single root CA issues certificates directly to users and computers. This is a simple model often used on private networks where there is no need for intermediate CAs.

In the single CA model, the root CA issues certificates to multiple intermediate CAs is incorrect because in the single CA model, the root CA issues certificates directly to users and computers, not to intermediate CAs.

The single CA model is often used in large, complex networks due to its scalability is incorrect because the single CA model is typically used in smaller, private networks due to its simplicity, not in large, complex networks.

The single CA model is highly secure and if the CA server is compromised, it does not affect the entire PKI is incorrect because if the single CA server is compromised in this model, the entire PKI is affected as all certificates are issued directly by this CA.

q_cert_concepts_third-party_ca_secp8

You are a security manager for a mid-sized company and are considering using a third-party certificate authority (CA) to manage your company's certificates.

Which of the following would be the MOST significant benefit of using a third-party CA?

Answers:

- It would allow your company to issue its own root certificates.
- It would eliminate the need for your company to manage any certificates.

- ***It would allow your company to set up different certificate policies through intermediate CAs.**
- It would make your company's certificates immune to revocation.

Explanation:

It would allow your company to set up different certificate policies through intermediate CAs is the correct answer. Third-party CAs often operate a hierarchical model where the root CA issues certificates to one or more intermediate CAs. This allows for different certificate policies, enabling clear understanding of what a particular certificate is designed for.

It would allow your company to issue its own root certificates is incorrect because when using a third-party CA, the CA issues the root certificates, not the company.

It would eliminate the need for your company to manage any certificates is incorrect because even when using a third-party CA, the company still needs to manage the certificates issued to its users and computers.

It would make your company's certificates immune to revocation is incorrect because all certificates, regardless of who issues them, can be revoked if necessary.

4.0 Identity and Access Management

4.1 Access Control Models

As you study this section, answer the following questions:

- What is access control and why is it important?
- How are rule-based access control and mandatory access control (MAC) similar?
- How does role-based control differ from rule-based control?
- How do separation of duties and job rotation differ?
- Which authentication type requires you to prove your identity?

The key terms for this section include:

Term	Definition
CIA Triad	Three principles of security control and management - confidentiality, integrity, and accessibility. Also known as the information security triad.
Confidentiality	The fundamental security goal of keeping information and communications private and protecting them from unauthorized access.
Integrity	The fundamental security goal of keeping organizational information accurate, free of errors, and without unauthorized modifications.
Availability	The fundamental security goal of ensuring that computer systems operate continuously and that authorized persons can access data that they need.
Non-repudiation	The security goal of ensuring that the party that sent a transmission or created data remains associated with that data and cannot deny sending or creating that data.
National Institute of Standards and Technology (NIST)	Develops computer security standards used by US federal agencies and publishes cybersecurity best practice guides and research.
cybersecurity frameworks (CSF)	Standards, best practices, and guidelines for effective security risk management. Some frameworks are general in nature, while others are specific to industry or technology types.
security controls	A technology or procedure put in place to mitigate vulnerabilities and risk and to ensure the confidentiality, integrity, and availability (CIA) of information.

Gap analysis	An analysis that measures the difference between the current and desired states in order to help assess the scope of work included in a project.
identity and access management (IAM)	A security process that provides identification, authentication, and authorization mechanisms for users, computers, and other entities to work with organizational assets like networks, operating systems, and applications.
Identification	The process by which a user account (and its credentials) is issued to the correct person. Sometimes referred to as enrollment.
Authentication	A method of validating a particular entity's or individual's unique credentials.
Authorization	The process of determining what rights and privileges a particular entity has.
Accounting	Tracking authorized usage of a resource or use of rights by a subject and alerting when unauthorized use is detected or attempted.
authentication, authorization, and accounting (AAA)	A security concept where a centralized platform verifies subject identification, ensures the subject is assigned relevant permissions, and then logs these actions to create an audit trail.
control plane	In zero trust architecture, functions that define policy and determine access decisions.
permissions	Security settings that control access to objects including file system items and network resources.
Discretionary access control (DAC)	An access control model where each resource is protected by an access control list (ACL) managed by the resource's owner (or owners).
Mandatory access control (MAC)	An access control model where resources are protected by inflexible, system-defined rules. Resources (objects) and users (subjects) are allocated a clearance level (or label).
Role-based access control (RBAC)	An access control model where resources are protected by ACLs that are managed by administrators and that provide user permissions based on job functions.
group account	A group account is a collection of user accounts that is useful when establishing file permissions and user rights because when many individuals need the same level of access, a group could be established containing all the relevant users.

Attribute-based access control (ABAC)	An access control technique that evaluates a set of attributes that each subject possesses to determine if access should be granted.
Rule-based access control	A nondiscretionary access control technique that is based on a set of operational rules or restrictions to enforce a least privileges permissions policy.
Least privilege	A basic principle of security stating that something should be allocated the minimum necessary rights, privileges, or information to perform its role.
Provisioning	The process of deploying an account, host, or application to a target production environment. This involves proving the identity or integrity of the resource, and issuing it with credentials and access permissions.
Deprovisioning	The process of removing an account, host, or application from the production environment. This requires revoking any privileged access that had been assigned to the object.
security identifier (SID)	The value assigned to an account by Windows and that is used by the operating system to identify that account.
group policy objects (GPOs)	On a Windows domain, a way to deploy per-user and per-computer settings such as password policy, account restrictions, firewall status, and so on.
geolocation	The identification or estimation of the physical location of an object, such as a radar source, mobile phone, or Internet-connected computing device.
time-of-day restrictions policy	Policies or configuration settings that limit a user's access to resources.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SYO-701	1.2 Summarize fundamental security concepts.
	• Confidentiality, Integrity, and Availability (CIA)
	• Non-repudiation
	• Authentication, Authorization, and Accounting (AAA)
	  ○ Authenticating people
	  ○ Authenticating systems
	  ○ Authorization models
	• Gap analysis
	• Zero trust

- Control plane
  - Adaptive identity
  - Threat scope reduction
  - Policy-driven access control
  - Policy Administrator
  - Policy Engine
- Data plane
  - Implicit trust zones
  - Subject/System
  - Policy enforcement point

2.5 Explain the purpose of mitigation techniques used to secure the enterprise.

- Access control
- Least privilege

4.6 Given a scenario, implement and maintain identity and access management.

- Provisioning/de-provisioning user accounts
- Permission assignments and implications
- Identity proofing
- Access controls
  - Mandatory
  - Discretionary
  - Role-based
  - Rule-based
  - Attribute-based
  - Time-of-day restrictions
  - Least privilege
- Multifactor authentication
  - Implementations
    - Hard/soft authentication tokens
  - Factors
    - Something you know
    - Something you have
    - Something you are
    - Somewhere you are

5.1 Summarize elements of effective security governance.

- Policies
  - Information security policies

4.1.1 Fundamental Security Concepts (Lesson Video)

Transcript:

In this lesson, we'll look at fundamental concepts that underpin the world of information security. Let's dive in.

Information security, often referred to as "infosec," is all about safeguarding data resources from unauthorized access, attacks, theft, or damage. Data can be vulnerable at rest, in transit, or during processing.

To ensure secure information, we focus on confidentiality, integrity, and availability. To balance these needs, we use the CIA Triad. First, Confidentiality. This means that information can only be accessed by authorized individuals. Second, Integrity ensures that data remains unaltered and is only modified by authorized parties. And third, Availability. This means information must be readily accessible to those authorized to use it. These three principles form the backbone of information security, and as security specialists, you'll find you're always seeking the best balance of all three components.

An access control system ensures that an information system meets the goals of the CIA triad. Access control dictates how subjects (like users, devices, or processes) interact with objects (like networks, servers, databases, apps, or files). Modern access control is implemented through Identity and Access Management, or IAM, consisting of four main processes: Identification—here, we create unique accounts for users, devices, or processes; Authentication—this is the process of proving the identity of subjects through credentials; Authorization—this determines and enforces rights for subjects on resources; and Accounting—this is the tracking of authorized resource usage and the detection of unauthorized access. Access control is essential for both people and systems, ensuring that only legitimate actions are permitted.

Some security models and researchers identify other properties of secure systems. The most important of these is non-repudiation. Non-repudiation means a person can't deny doing something, such as creating, modifying, or sending a resource.

For example, providing proof of authenticity and integrity of data that's sent or received.

Gap analysis is a process that assesses how well an organization's security systems align with a cybersecurity framework. Security functions have specific goals achieved through security controls. To select the right controls, organizations use frameworks to guide them.

Frameworks provide structure and guidance for security programs. Gap analysis identifies deviations from framework requirements and recommends remediation, thus helping organizations maintain cybersecurity, meet compliance requirements, and prioritize investments.

Zero Trust is a security model that treats all devices, users, and services as untrusted, regardless of location within or outside the network's perimeter.

This model relies on fundamental concepts for comprehensive security. These include adaptive identity, threat score reduction, and policy-driven access control. Adaptive identity recognizes that identities change and need continuous verification. Threat scope reduction limits access to resources on a need-to-know basis, reducing attack surfaces. Policy-driven access control enforces access restrictions based on user identity, device posture, and network context.

Zero Trust separates the control plane and data plane for flexibility and scalability. In a Zero Trust architecture, the control plane manages policies and authorization, while the data plane handles secure data transfer.

The control plane defines access policies, monitors for threats, and updates policies. The data plane establishes secure data pathways and enforces access policies. Separating these planes enhances flexibility and security, allowing for granular control and continuous monitoring.

And that's it for this lesson. In this lesson, we introduced you to fundamental information security concepts. We first looked at information security applied through the CIA Triad. Then, we discussed access control, which establishes how subjects interact with objects in a system. Next, we talked about gap analysis and how it helps you know where to improve your security programs. We finished this lesson by reviewing Zero Trust security concepts and how those concepts are applied through control planes and data planes in a Zero Trust model.

4.1.2 Access Control Facts

This lesson covers the following topics:

- Information security
- Cybersecurity framework
- Gap analysis
- Access control
- Zero Trust security concepts
- Control and data planes in Zero Trust models

Information Security

Information security (infosec) refers to the protection of data resources from unauthorized access, attack, theft, or damage. Data may be vulnerable because of the way it is stored, transferred, or processed. The systems used to store, transmit, and process data must demonstrate the properties of security. Secure information has three properties, often referred to as the CIA Triad :

- **Confidentiality** means that information can only be read by people who have been explicitly authorized to access it.
- **Integrity** means that the data is stored and transferred as intended and that any modification is authorized.
- **Availability** means that information is readily accessible to those authorized to view or modify it.

The triad can also be referred to as "AIC" to avoid confusion with the Central Intelligence Agency.

Some security models and researchers identify other properties of secure systems. The most important of these is non-repudiation. Non-repudiation means that a person cannot deny doing something, such as creating, modifying, or sending a resource. For example, a legal document, such as a will, must usually be witnessed when it is signed. If there is a dispute about whether the document was correctly executed, the witness can provide evidence that it was.

Cybersecurity Framework

Within the goal of ensuring information security, cybersecurity refers specifically to provisioning secure processing hardware and software. Information security and cybersecurity tasks can be classified into five functions, following the framework developed by the National Institute of Standards and Technology (NIST) (nist.gov/cyberframework/online-learning/five-functions):

- **Identify** — develop security policies and capabilities. Evaluate risks, threats, and vulnerabilities and recommend security controls to mitigate them.
- **Protect** — procure/develop, install, operate, and decommission IT hardware and software assets with security as an embedded requirement of every stage of this operation's lifecycle.
- **Detect** — perform ongoing, proactive monitoring to ensure that controls are effective and capable of protecting against new types of threats.
- **Respond** — identify, analyze, contain, and eradicate threats to systems and data security.
- **Recover** — implement cybersecurity resilience to restore systems and data if other controls are unable to prevent attacks.

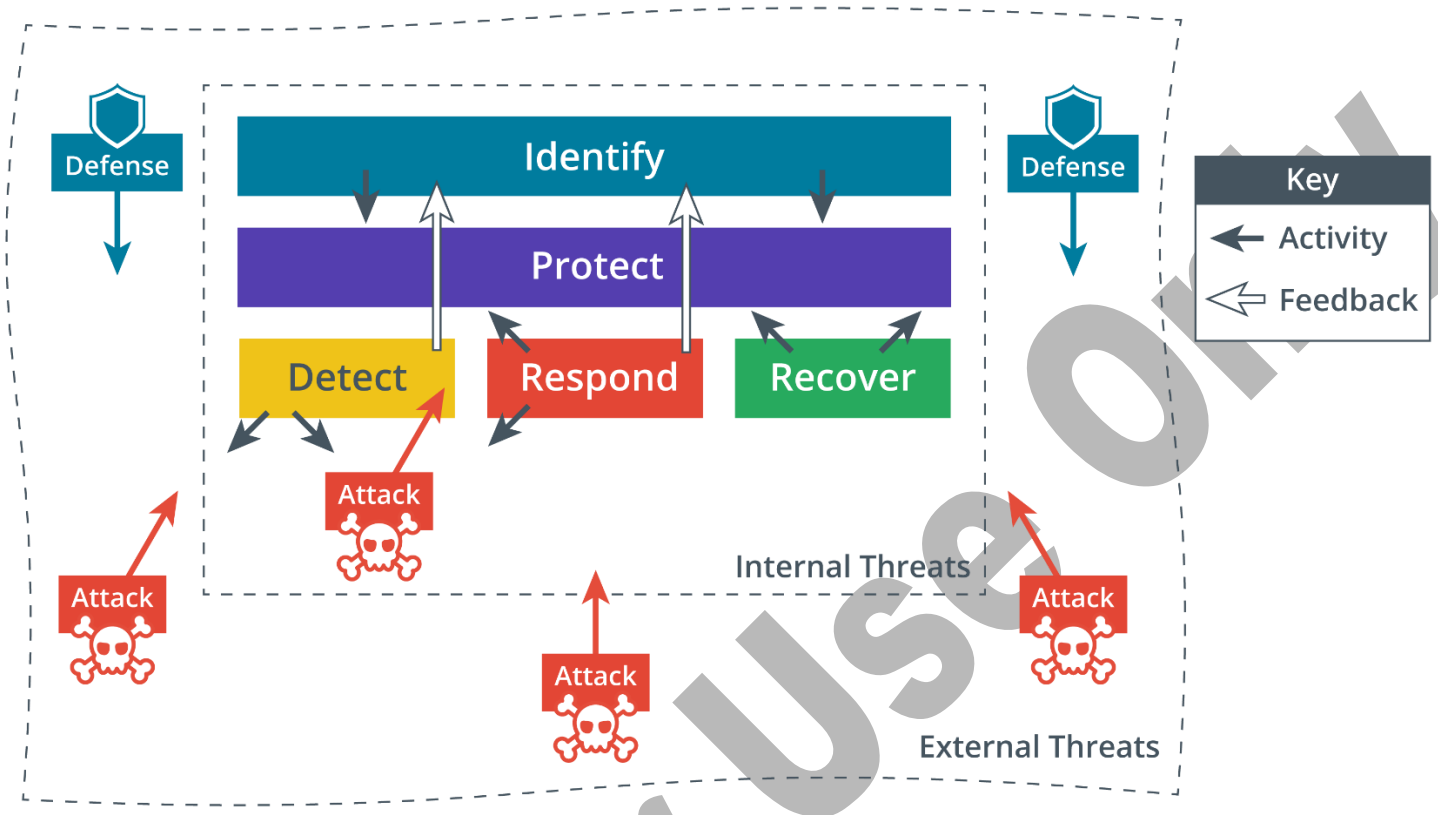

Core cybersecurity tasks.

NIST's framework is just one example. There are many other cybersecurity frameworks (CSF) .

Gap Analysis

Each security function is associated with a number of goals or outcomes. For example, one outcome of the Identify Function is an inventory of the assets owned and operated by the company. Outcomes are achieved by implementing one or more security controls .

Numerous categories and types of security controls cover a huge range of functions. This makes the selection of appropriate and effective controls difficult.

A cybersecurity framework guides the selection and configuration of controls. Frameworks are important because they save an organization from building its security program in a vacuum or building it on a foundation that fails to account for important security concepts.

The use of a framework allows an organization to make an objective statement of its current cybersecurity capabilities, identify a target level of capability, and prioritize investments to achieve that target. This gives a structure to internal risk management procedures and provides an externally verifiable statement of regulatory compliance.

Gap analysis is a process that identifies how an organization's security systems deviate from those required or recommended by a framework. This will be performed when first adopting a framework or when meeting a new industry or legal compliance requirement. The analysis might be repeated every few years to meet compliance requirements or to validate any changes that have been made to the framework.

For each section of the framework, a gap analysis report will provide an overall score, a detailed list of missing or poorly configured controls associated with that section, and recommendations for remediation.

Function	Controls (Actual/Required)	CIA Triad Risk Levels	Target Remediation
Identify (10/16)	Asset Management (4/6)	C : 6 I : 6 A : 6	Q4
	Governance (3/4)	C : 6 I : 6 A : 1	Q3
	Risk Assessment (3/6)	C : 6 I : 6 A : 3	Q3
Protect (8/16)	Identity and Access Management (5/8)	C : 9 I : 9 A : 4	Q1
	Data Security (3/8)	C : 9 I : 9 A : 4	Q1

Advanced capability
Intermediate capability
No/basic capability

Summary of gap analysis findings showing the number of recommended controls not implemented per function and category, plus risks to confidentiality, integrity, and availability from missing controls and target remediation date.

While some or all work involved in gap analysis could be performed by the internal security team, a gap analysis is likely to involve third-party consultants. Frameworks and compliance requirements from regulations and legislation can be complex enough to require a specialist. Advice and feedback from an external party can alert the internal security team to oversights, new trends, and changes in best practices.

Access Control

An access control system ensures that an information system meets the goals of the CIA triad. Access control governs how subjects/principals may interact with objects. Subjects are people, devices, software processes, or any other system that can request and be granted access to a resource. Objects are the resources. An object could be a network, server, database, app, or file. Subjects are assigned rights or permissions on resources.

Modern access control is typically implemented as an identity and access management (IAM) system. IAM comprises four main processes:

- **Identification** — creating an account or ID that uniquely represents the user, device, or process on the network.
- **Authentication** — proving that a subject is who or what it claims to be when it attempts to access the resource. An authentication factor determines what sort of credential the subject can use. For example, people might be authenticated by providing a password; a computer system could be authenticated using a token such as a digital certificate.
- **Authorization** — determining what rights subjects should have on each resource and enforcing those rights. An authorization model determines how these rights are granted. For example, in a discretionary model, the object owner can allocate rights. In a mandatory model, rights are predetermined by system-enforced rules and cannot be changed by any user within the system.
- **Accounting** — tracking authorized usage of a resource or use of rights by a subject and alerting when unauthorized use is detected or attempted.

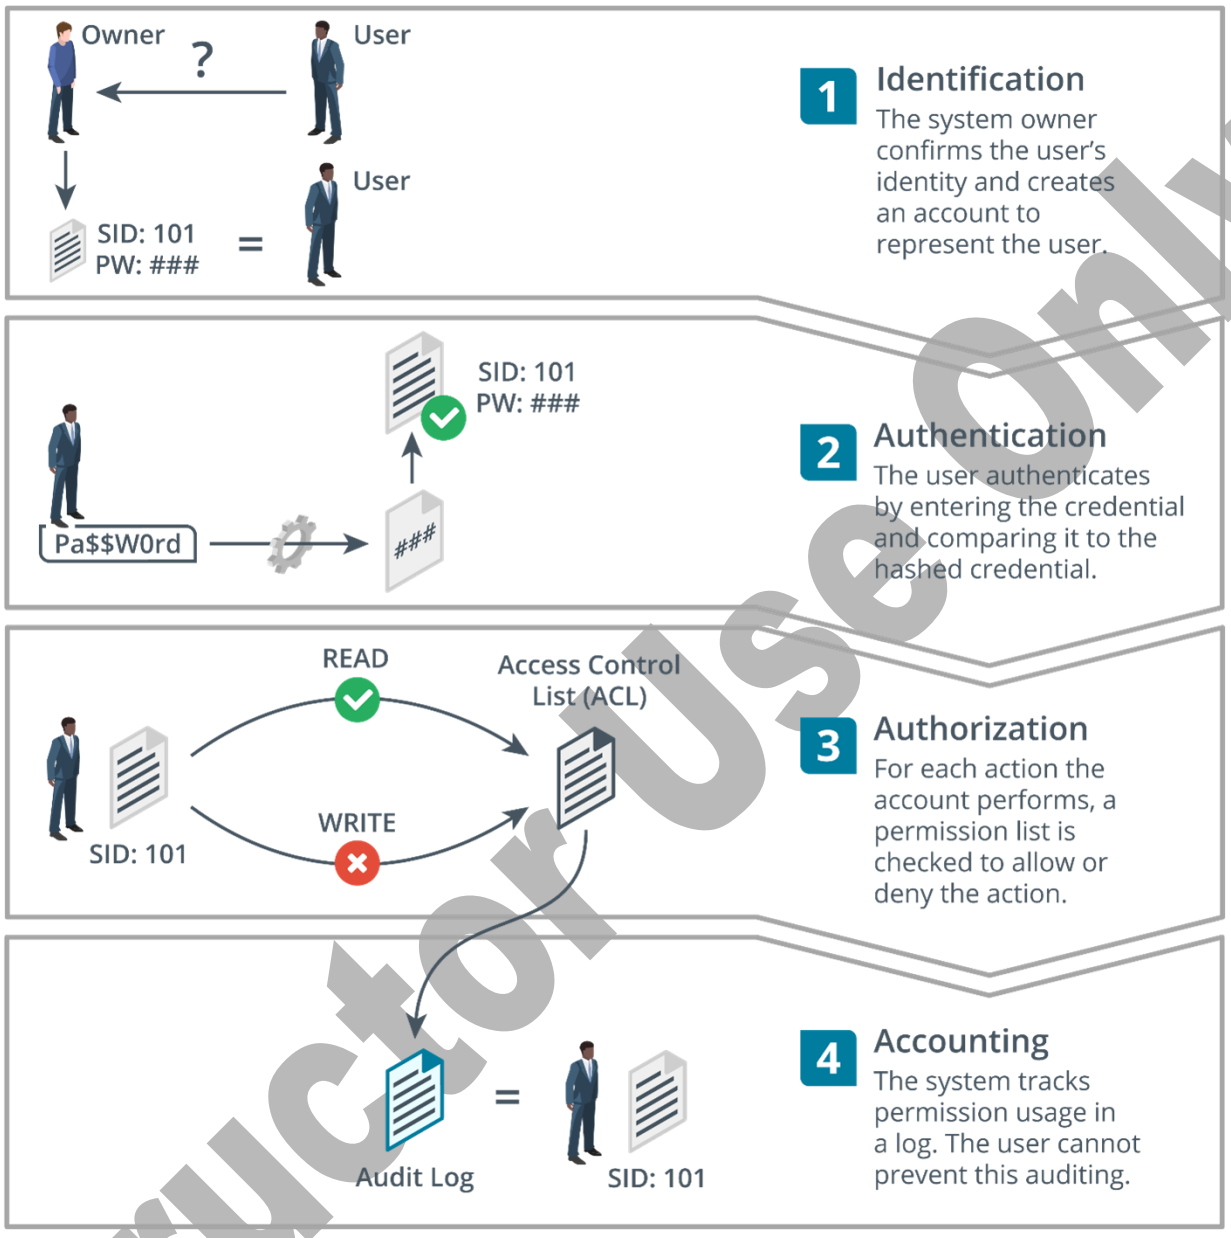

Differences among identification, authentication, authorization, and accounting. (Images © 123RF.com.)

The servers and protocols that implement these functions can also be referred to as authentication, authorization, and accounting (AAA) . The use of IAM to describe enterprise security workflows is becoming more prevalent as the importance of the identification process is better acknowledged.

For example, if you are setting up an e-commerce site and want to enroll users, you need to select the appropriate controls to perform each function:

- **Identification** — ensure that customers are legitimate. For example, you might need to ensure that billing and delivery addresses match and that they are not trying to use fraudulent payment methods.
- **Authentication** — ensure that customers have unique accounts and that only they can manage their orders and billing information.

- **Authorization** — rules to ensure customers can place orders only when they have valid payment mechanisms in place. You might operate loyalty schemes or promotions that authorize certain customers to view unique offers or content.
- **Accounting** — the system must record the actions a customer takes (to ensure that they cannot deny placing an order, for instance).

Remember that these processes apply both to people and to systems. For example, you need to ensure that your e-commerce server can authenticate its identity when customers connect to it using a web browser.

Zero Trust Security Concepts

Zero Trust is a security model that assumes that all devices, users, and services are not inherently trusted, regardless of whether inside or outside a network's perimeter. Instead, the Zero Trust model requires all users and devices to be authenticated and authorized before accessing network resources. The Zero Trust model includes several fundamental concepts that provide a comprehensive security solution.

- **Adaptive identity** recognizes that user identities are not static and that identity verification must be continuous and based on a user's current context and the resources they are attempting to access.
- **Threat scope reduction** means that access to network resources is granted on a need-to-know basis, and access is limited to only those resources required to complete a specific task. This concept reduces the network's attack surface and limits the damage a successful attack can cause.
- **Policy-driven access control** describes how access control policies are used to enforce access restrictions based on user identity, device posture, and network context.

Device posture refers to the security status of a device, including its security configurations, software versions, and patch levels. In a security context, device posture assessment involves evaluating the security status of a device to determine whether it meets certain security requirements or poses a risk to the network.

Control and Data Planes in Zero Trust Models

In a Zero Trust architecture, the control and data planes are implemented separately and have different functions.

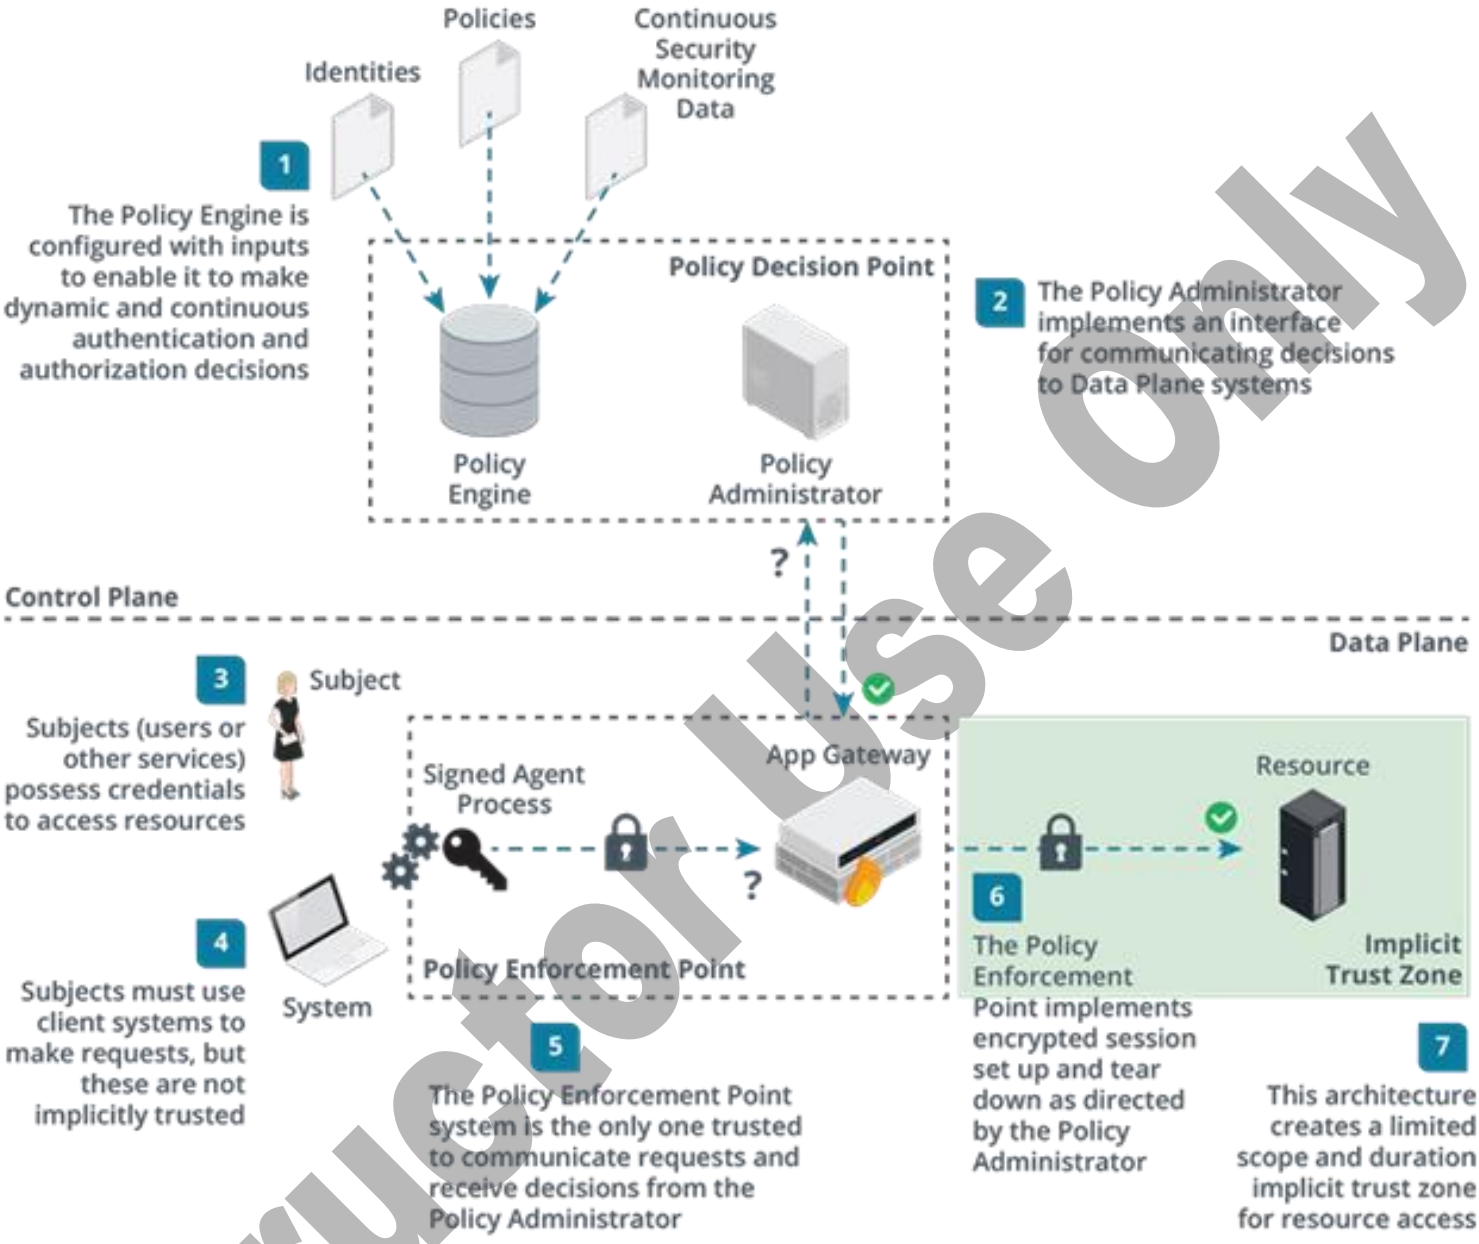

Components in NIST's Zero Trust architecture framework.

The control plane manages policies that dictate how users and devices are authorized to access network resources. It is implemented through a centralized policy decision point. The policy decision point is responsible for defining policies that limit access to resources on a need-to-know basis, monitoring network activity for suspicious behavior, and updating policies to reflect changing network conditions and security threats. The policy decision point is comprised of two subsystems:

- The policy engine is configured with subject and host identities and credentials, access control policies, up-to-date threat intelligence, behavioral analytics, and other results of host and network security scanning and monitoring. This comprehensive state data allows it to define an algorithm and metrics for making dynamic authentication and authorization decisions on a per-request basis.
- The policy administrator is responsible for managing the process of issuing access tokens and establishing or tearing down sessions based on the decisions made by the policy engine. The policy administrator implements an interface between the control and the data plane.

Where systems in the control plane define policies and make decisions, systems in the data plane establish sessions for secure information transfers. In the data plane, a subject (user or service) uses a system (such as a client host PC, laptop, or smartphone) to make requests for a given resource. A resource is typically an enterprise app running on a server or cloud. Each request is mediated by a policy enforcement point. The enforcement point might be implemented as a software agent running on the client host that communicates with an app gateway. The policy enforcement point interfaces with the policy administrator to set up a secure data pathway if access is approved or tear down a session if access is denied or revoked.

The processes implementing the policy enforcement point are the only ones permitted to interface with the policy administrator. It is critical to establish a root of trust for these processes so that policy decisions cannot be tampered with.

The data pathway established between the policy enforcement point and the resource is referred to as an implicit trust zone. For example, the outcome of a successful access request might be an IPsec tunnel established between a digitally signed agent process running on the client, a trusted web application gateway, and the resource server. Because the data is protected by IPsec transport encryption, no tampering by anyone with access to the underlying network infrastructure (switches, access points, routers, and firewalls) is possible.

The goal of Zero Trust design is to make this implicit trust zone as small as possible and as transient as possible. Trusted sessions might only be established for individual transactions. This granular or microsegmented approach is in contrast with perimeter-based models, where trust is assumed once a user has authenticated and joined the network. In Zero Trust, place in the network is not a sufficient reason to trust a subject request. Similarly, even if a user is nominally authenticated, behavioral analytics might cause a request to be blocked or a session to be terminated.

Separating the control and data plane is significant because it allows for a more flexible and scalable network architecture. The centralized control plane ensures consistency for access request handling across both the managed enterprise network and unmanaged internet or third-party networks, regardless of the devices being used or the user's location. This makes managing access control policies and monitoring network activity for suspicious behavior easier. Continuous monitoring via the independent control plane means that sessions can be terminated if anomalous behavior is detected.

4.1.3 Access Control Best Practices

Access control refers to regulating and managing the permissions granted to individuals, software, systems, and networks to access resources or information. Access controls ensure that only authorized entities can perform specific actions or access certain data, while unauthorized entities are denied access. Access control concepts apply to networks, physical access, data, applications, and the cloud.

This lesson covers the following topics:

- Access control best practices
- Transition best practices

Access Control Best Practices

Access control best practices take into consideration the following security principles and concepts:

Principle	Description
Principle of least privilege	Implementing the principle of least privilege (PoLP) is a cornerstone of improving endpoint protection and minimizing the risk of security issues. The principle of least privilege dictates that users, applications, and processes should only be granted the minimum permissions necessary to complete their duties and nothing more.

Principle	Description
	There are several practical methods for implementing least privilege. An essential first step to effectively implementing least privilege is thoroughly auditing user roles, privileges, and responsibilities. This process allows organizations to understand what access each user needs to perform their job role effectively. Access controls and permissions can be adjusted to adopt a principle of least privilege that best reflects the audit results. User and account management tools are also essential when implementing the principle of least privilege. Regularly reviewing and removing unused or unnecessary accounts reduces the potential targets for an attacker. Similarly, temporary privileges, which grant additional access rights for a limited time and only when required, can help keep privileges as restrictive as possible. The principle of least privilege also applies to software applications and operating systems, not just to users. For instance, ensuring that applications run with the minimum necessary permissions can prevent them from being exploited to carry out privileged actions. Common methods of controlling access include: • *Implicit deny* denies access to users or groups not explicitly given access to a resource. Implicit deny is the weakest form of privilege control. • *Explicit allow* specifically identifies users or groups who have access. Explicit allow is a moderate form of access control in which privilege has been granted to a subject. • *Explicit deny* identifies users or groups who are not allowed access. Explicit deny is the most robust access control, overruling all other privileges granted.
Need to know	Need to know describes the restriction of highly sensitive data usually referenced in government and military contexts. Essential facts about the need to know include: • Even if an individual is fully cleared, the information will not be divulged unless the person needs to know the information to perform official duties. • Need to know discourages casual browsing of sensitive materials. • In a classified environment, a clearance into a top-secret compartment allows access to only specific information. This is a form of mandatory access control (MAC).
Separation of duties	Separation of duties is the concept of having more than one person required to complete a task. This is a preventive principle primarily designed to reduce conflicts of interest. It also prevents insider attacks because no one person has end-to-end control, and no one person is irreplaceable. Essential facts to know about separation of duties include: • A business can use the principle of split knowledge to achieve a separation of duties. This means no single person controls a system's security mechanisms; no single person can completely compromise the system. • In sensitive or high-risk transactions, a business can use two-man controls. This means that two operators must review and approve each other's work.
Job rotation	Job rotation is a technique where users are cross-trained in multiple positions. Responsibilities are regularly rotated between personnel. Job rotation:

Principle	Description
	• Cross trains staff in different functional areas to detect fraud. • Exchanges positions of two or more employees to allow for oversight of past transactions. • Can be used for training purposes.
Defense-in-depth	Defense-in-depth is an access control principle that implements multiple access control methods instead of relying on a single process. Numerous defenses make it harder to bypass security measures.
Identification	Identification is claiming an identity, such as telling someone your name. Essential facts to know about identification include: • A username is a form of identification. • Identification is not very secure because anyone could pretend to be the user. • To substantiate identity, the person must provide some form of identity verification.
Multi-Factor Authentication	Multi-Factor Authentication is using more than one way to verify identity. Multi-Factor Authentication is achieved by requiring two or more methods that only the user can provide. Five categories of computer system authentication include: • *Something you are* , such as biometric information (e.g., fingerprint or retina scan). • *Something you have* , such as smart cards, RSA tokens, or security key fobs. • *Something you know* , such as passwords and PINs. • *Somewhere you are* , such as a geographical location. • *Something you do* , such as how you type a sentence on a keyboard.
Mutual authentication	Mutual authentication is when two communicating entities authenticate each other before exchanging data. It requires not only the server to authenticate the user but the user to authenticate the server. This makes mutual authentication more secure than traditional, one-way authentication.
Time of day restrictions	Time of day restrictions impose restrictions on incoming and outgoing network traffic based on the time of day. This allows organizations to define when resources can be accessed or specific actions can be performed. This can help to: • Limit access to critical systems after hours. • Prevent unauthorized activities during particular periods. • Reduce the attack surface available to potential threats.

Transition Best Practices

Organizations should follow strict guidelines when an employee transitions from or into a new position. Creeping privileges occur when a user's job position changes, and the user is granted a new set of access privileges. Still, the user's current access privileges are not removed or modified, resulting in privilege escalation. As a result, the user accumulates unnecessary privileges for the current work tasks. The principle of least privilege and separation of duties are

countermeasures against creeping privileges. To avoid creeping privileges and to best protect the security of information, the following precautions should be taken in each stage of the account's life cycle:

Event	Precautions
Account creation	When creating an account, apply the appropriate access rights based on the job role implemented in the access control system. Use the principle of least privilege and grant only the minimum privileges required to perform the position's duties.
Active accounts	During the life of an account: • Modify access rights as job roles and circumstances change. • Monitor password resets and lockouts to ensure account security. • Re-evaluate access rights periodically.
Old accounts	When an account is no longer needed, take appropriate actions to: • Delete accounts that will no longer be used. • Rename accounts to give new users in the same job role the same access privileges. • Lock accounts that will not be used for extended periods to prevent them from being used. • Remove unnecessary rights from accounts that will be kept on the system. • Archive important data or files the user owns or assign ownership to another user. • Prohibit generic user accounts, such as the Guest or Administrator users, on Windows systems. • End-of-life procedures should include deactivating or deleting unused accounts and destroying data that might remain on storage media. This will prevent sensitive data from being accessible to unauthorized users.

4.1.4 Access Control Models (Lesson Video)

Transcript:

Depending on your organization's needs, there are several ways to implement access control. In this lesson, we'll look at some common implementation models.

Let's start with the least restrictive model, Discretionary Access Control, or DAC. With DAC, every resource has an owner. That owner decides who has access to the resource. DAC uses programs like access control lists, or ACLs. A good example is NTFS within a Windows system.
Each NTFS file and folder has an owner. The owner can go into the file or folder's access control list and decide which user or group can access the resource. Because access is resource-specific, a lot of administrative work is required to implement and maintain these policies on each individual resource.
Mandatory access control, or MAC, is a static system with classification labels and levels. When you're using MAC, every computer, every file, and every object is assigned a label. These labels indicate how important the object is. Each user is assigned an access level. MAC compares the object labels with the user's level to grant or deny access to a given resource.
Let's say your organization uses three labels: Confidential, Secret, and Top Secret. Your organization also uses three user access levels. Users with access level 1 can see anything labeled Confidential. They cannot see anything labeled

Secret or Top Secret. Users with access level 2 can see anything labeled Confidential or Secret, but cannot see anything labeled Top Secret. Users with access level 3 can see anything labeled Confidential, Secret, or Top Secret.
It's important to note that permissions can't be altered for specific instances. For example, the data owner can't prevent specific people with level 2 clearance from accessing certain documents labeled as level 1. It's an all-or-none deal.
Next, we have role-based access control. This model is a hybrid between MAC and DAC, and it's probably the most commonly used. Role-based access control means that your role within the organization determines whether or not you're able to access certain kinds of data.
For example, Mary is the CFO, so she needs full access to important company financial records, while Aubrey is the receptionist and will only need to see calendar information for everyone in the organization.
We also have rule-based access controlâ€""don't confuse it with role-based! Rule-based access control is used with routers through router access control lists. With a router access control list, or ACL, you decide which IPs are allowed through the router using rules. Either you match that rule, or you don't. This has nothing to do with your user account, what group you're a member of, or even a classification label. The only thing that matters is whether the condition meets the rules configured on the router.
Attribute-Based Access Control, or ABAC, provides more flexibility because it combines object attributes to determine access. These attributes can be a number of different things, such as a user's role, position, or even project association. ABAC rules use an if-then-else format. If the user has the required attribute, they're granted access. If the user doesn't have the required attribute, they're denied access.
For example, let's say we have a user named RBradbury. This user's role is manager. Their department is development, and they're currently working on a project called AmazingApp.
Now, let's say RBradbury wants to access a file on the server. This file has been configured with several attributes: role = manager, department = development, and project = BoringApp.
When RBradbury tries to access the file, the system asks, "Are you a manager?" Yes. "Is your department development?" Yes. "Are you on the BoringApp project?" No. Since it hit a false return, the user is denied access to the file.
Modern security concerns extend beyond a desktop in the office. Conditional Access is a way to enforce access control while also encouraging users to be productive wherever they are. Conditional access isn't intended to be the first point of security. Instead, it steps in after first-factor authentication has been granted. Conditional access can be configured to consider many different factors.
Administrators can maintain specific control at the user or group level. They can permit or deny access based on an IP address or an IP range. Administrators can also use specific applications or devices to permit, restrict, or deny access to users.
That's it for this lesson. In this video, we discussed several access control models: discretionary access control, mandatory access control, role-based access control, rule-based access control, attribute-based access control, and conditional access control.

4.1.5 Access Control Model Facts

Authorization is the part of identity and access management that assigns privileges to network users and services. Implementing an access control model helps an organization manage the implications of privilege assignments and accounts for the actions of both regular and privileged administrative users. Account policies help you to protect credentials and to detect and manage risks from compromised accounts.

This lesson covers the following topics:

- Discretionary and mandatory access control
- Role- and attribute-based access control
- Rule-based access control
- Least privilege permission assignments
- Account attributes and access policies
- Location-based policies
- Time-Based Restrictions

Discretionary and Mandatory Access Control

A user account that has been authenticated can be allocated rights and permissions on networks, computers, and data. An access control model describes the principles that govern how users receive rights.

Discretionary Access Control

Discretionary access control (DAC) is based on the primacy of the resource owner. In a DAC model, every resource has an owner. The owner creates a file or service, although ownership can be assigned to another user. The owner has full control over the resource, and they can modify its access control list (ACL) to grant rights to others.

DAC is the most flexible model and is currently implemented widely in computer and network security. In file system security, it is the model used by default for most UNIX/Linux distributions and Microsoft Windows. As the most flexible model, it is also the weakest because it makes centralized administration of security policies the most difficult to enforce. It is also the easiest to compromise, as it is vulnerable to insider threats and abuse of compromised accounts.

Mandatory Access Control

The DAC model exposes information to the threat of compromise via the privileged owner accounts. Mandatory access control (MAC) is based on security clearance levels. Rather than defining ACLs on resources, each object is given a classification label, and each subject is granted a clearance level. In a confidentiality-oriented multilevel system, subjects are permitted to read objects classified at their own clearance level or below. For example, a user with Top-Secret clearance could read data with Top-Secret, Secret, and Confidential classification labels. A user with Secret clearance could access Secret and Confidential levels only.

Labeling objects and granting clearance takes place using pre-established rules. The critical point is that these rules are nondiscretionary and cannot be changed by any subject account.

As a simple classification system is inflexible, most MAC models add the concept of compartment-based access. For example, a data file might be at Secret classification and located in the HR compartment. Only subjects with Secret and HR clearance could access the file.

In MAC, users with high clearance are not permitted to write low-clearance documents. This is referred to as write up, read down. This prevents, for example, a user with Top Secret clearance from republishing some Top Secret data that they can access with Secret clearance.

Role- and Attribute-Based Access Control

Role- and attribute-based access control use nondiscretionary, rules-based permissions assignments with more flexibility than MAC.

Role-Based Access Control

Role-based access control (RBAC) means that an organization defines its permission requirements in terms of the tasks that an employee or service must be able to perform. Each set of permissions is a role. Each principal (user or service account) is allocated to one or more roles. Under this system, the right to modify the permissions assigned to each role is reserved to a system owner. Therefore, the system is nondiscretionary as each principal cannot modify the ACL of a resource, even though they can change the resource in other ways. Principals gain rights implicitly (through being assigned to a role) rather than explicitly (being assigned the right directly).

The concept of a security group account goes some way toward turning a discretionary system into a role-based one. Rather than assigning rights directly to user accounts, the system owner assigns user accounts to security group accounts. Principals gain rights by being made a member of a security group. A principal can be a member of multiple groups and, therefore, receive rights and permissions from several sources.

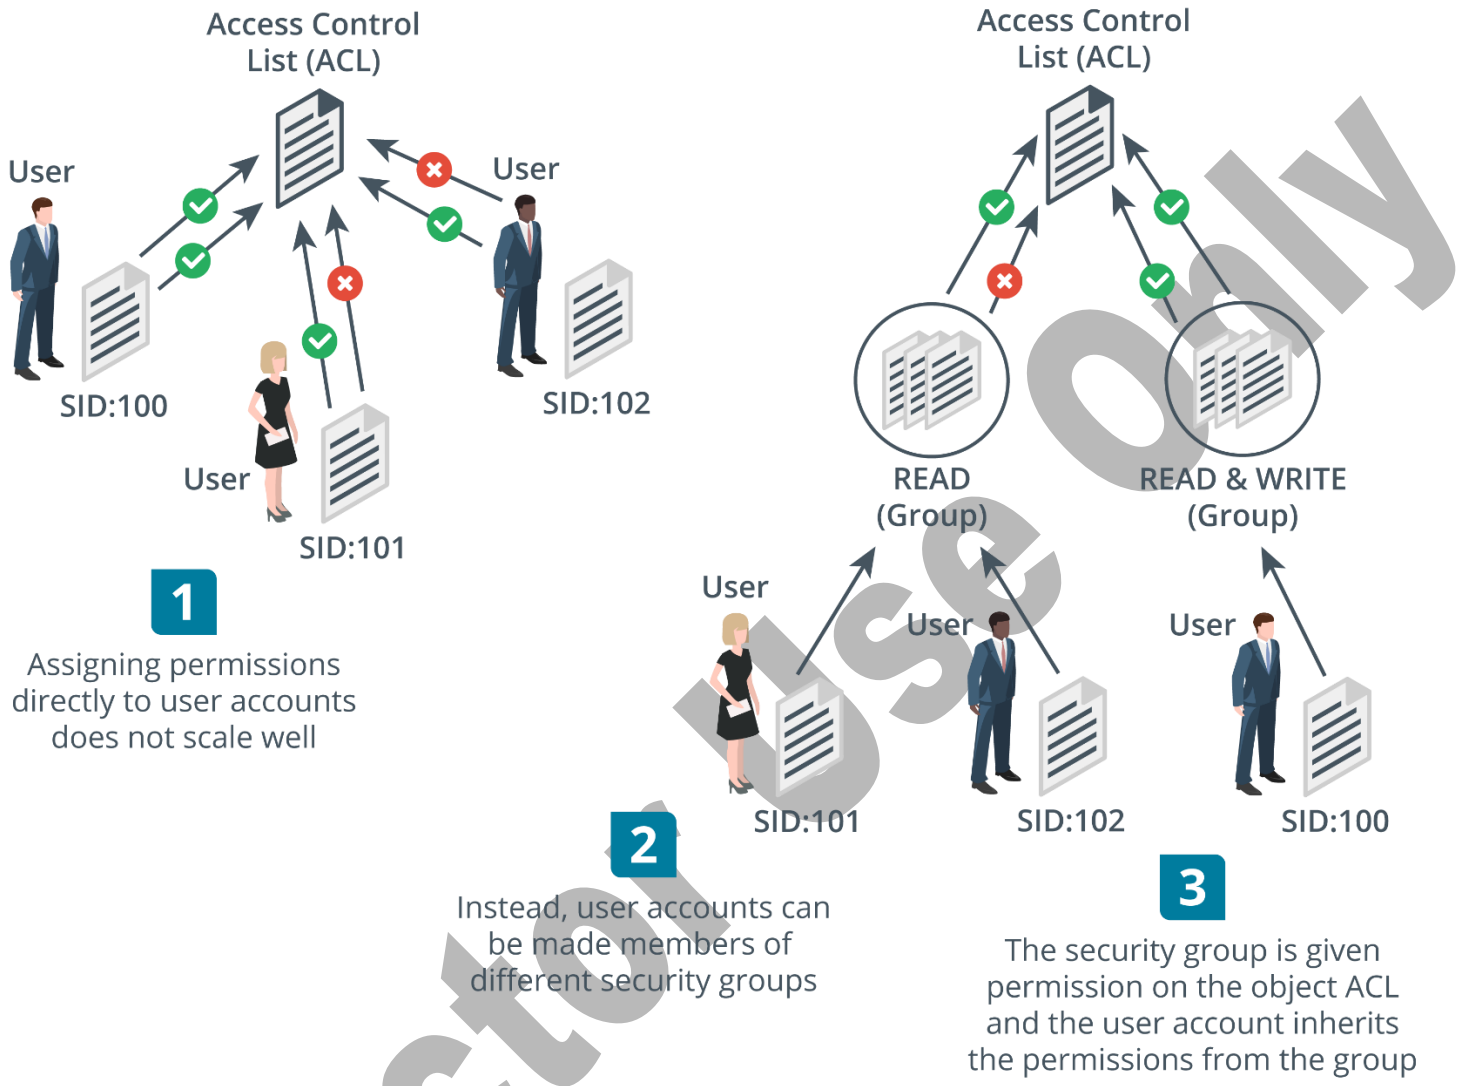

Using security groups to assign privileges. (Images © 123RF.com.)

RBAC can be partially implemented by mapping security groups onto roles, but they are not identical schemes. Membership in security groups is largely discretionary (assigned by administrators rather than determined by the system). Also, ideally, a principal should only inherit the permissions of a role to complete a particular task rather than retain them permanently. Administrators should be prevented from escalating their own privileges by assigning roles to their own accounts arbitrarily or boosting a role's permissions.

Attribute-Based Access Control

Attribute-based access control (ABAC) is the most fine-grained type of access control model. As the name suggests, an ABAC system makes access decisions based on a combination of subject and object attributes plus any context-sensitive or system-wide attributes. As well as group/role memberships, these attributes could include information about the OS currently being used, the IP address, or the presence of up-to-date patches and antimalware. An attribute-based system monitors the number of events or alerts associated with a user account or with a resource or tracks access requests to ensure they are consistent in terms of timing or geographic location. It can be programmed to implement policies such as M-of-N control and separation of duties.

The attributes assigned to a resource constitute a policy that uses Boolean logic to determine who can access the resource. An example of a file access policy might include the following attributes: role = manager, department = development, and project = NewApp. Only users who possess all three attributes can access the file.

Rule-Based Access Control

Rule-based access control refers to any sort of access control model where access control policies are determined by system-enforced rules rather than system users. As such, RBAC, ABAC, and MAC are all examples of rule-based (or nondiscretionary) access control.

Conditional access is an example of rule-based access control. A conditional access system monitors account or device behavior throughout a session. If certain conditions are met, it may suspend the account or require the user to reauthenticate, perhaps using a two-step verification method.

Conditional access is a way to enforce access control while encouraging users to be productive wherever they are. Conditional access isn't intended to be the first point of security. Instead, it steps in after the first-factor authentication has been granted. Conditional access policies ask a user to complete an action to access a resource. Depending on the level of security of the requested resource, the user may be required to complete more actions. For policy decisions, conditional access can be configured to consider many different factors, including:

- Implement control at the user or group level.
- Permit or deny access based on an IP address or an IP range.
- Permit or deny access to users who are using specific applications.
- Permit, restrict, or deny access to users using specific devices or device states.

The User Account Control (UAC) and sudo restrictions on privileged accounts are examples of conditional access. The user is prompted for confirmation or authentication when making requests that require elevated privileges. Role-based rights management and ABAC systems can apply a number of criteria to conditional access, including location-based policies.

An example of a rule-based access control implementation is a router access control list that allows or denies traffic based on characteristics within the packet, such as IP address or port number. Because rule-based access control does not consider the subject's identity, a system that uses rules can be regarded as mandatory access control.

Least Privilege Permission Assignments

Least privilege means that a principal is granted the minimum possible sufficient rights to complete a task they are authorized to perform. This mitigates risk if an account should be compromised and fall under the control of a threat actor. Least privilege involves a design phase, where analysis of business workflows determines what roles and permissions are required.

While least privilege is a strong design principle, successfully implementing it can be challenging. Where many users, groups, roles, and resources are involved, managing permission assignments and implications is complex and time-consuming. Improperly configured accounts can have two different impacts. On the one hand, setting privileges that are too restrictive creates a large volume of support calls and reduces productivity. On the other hand, granting too many privileges to users weakens the system's security and increases the risk of malware infection and a data breach.

Ensuring least privilege also involves continual monitoring to prevent authorization creep. Authorization creep refers to a situation where a user acquires more rights, either directly or by being added to security groups or roles.

For example, a user may be granted elevated privileges temporarily. In this case, a system is needed to ensure that the privileges are revoked at the end of the agreed period. A system of auditing should regularly review privileges, monitor group membership, review access control lists for each resource, and identify and disable unnecessary accounts.

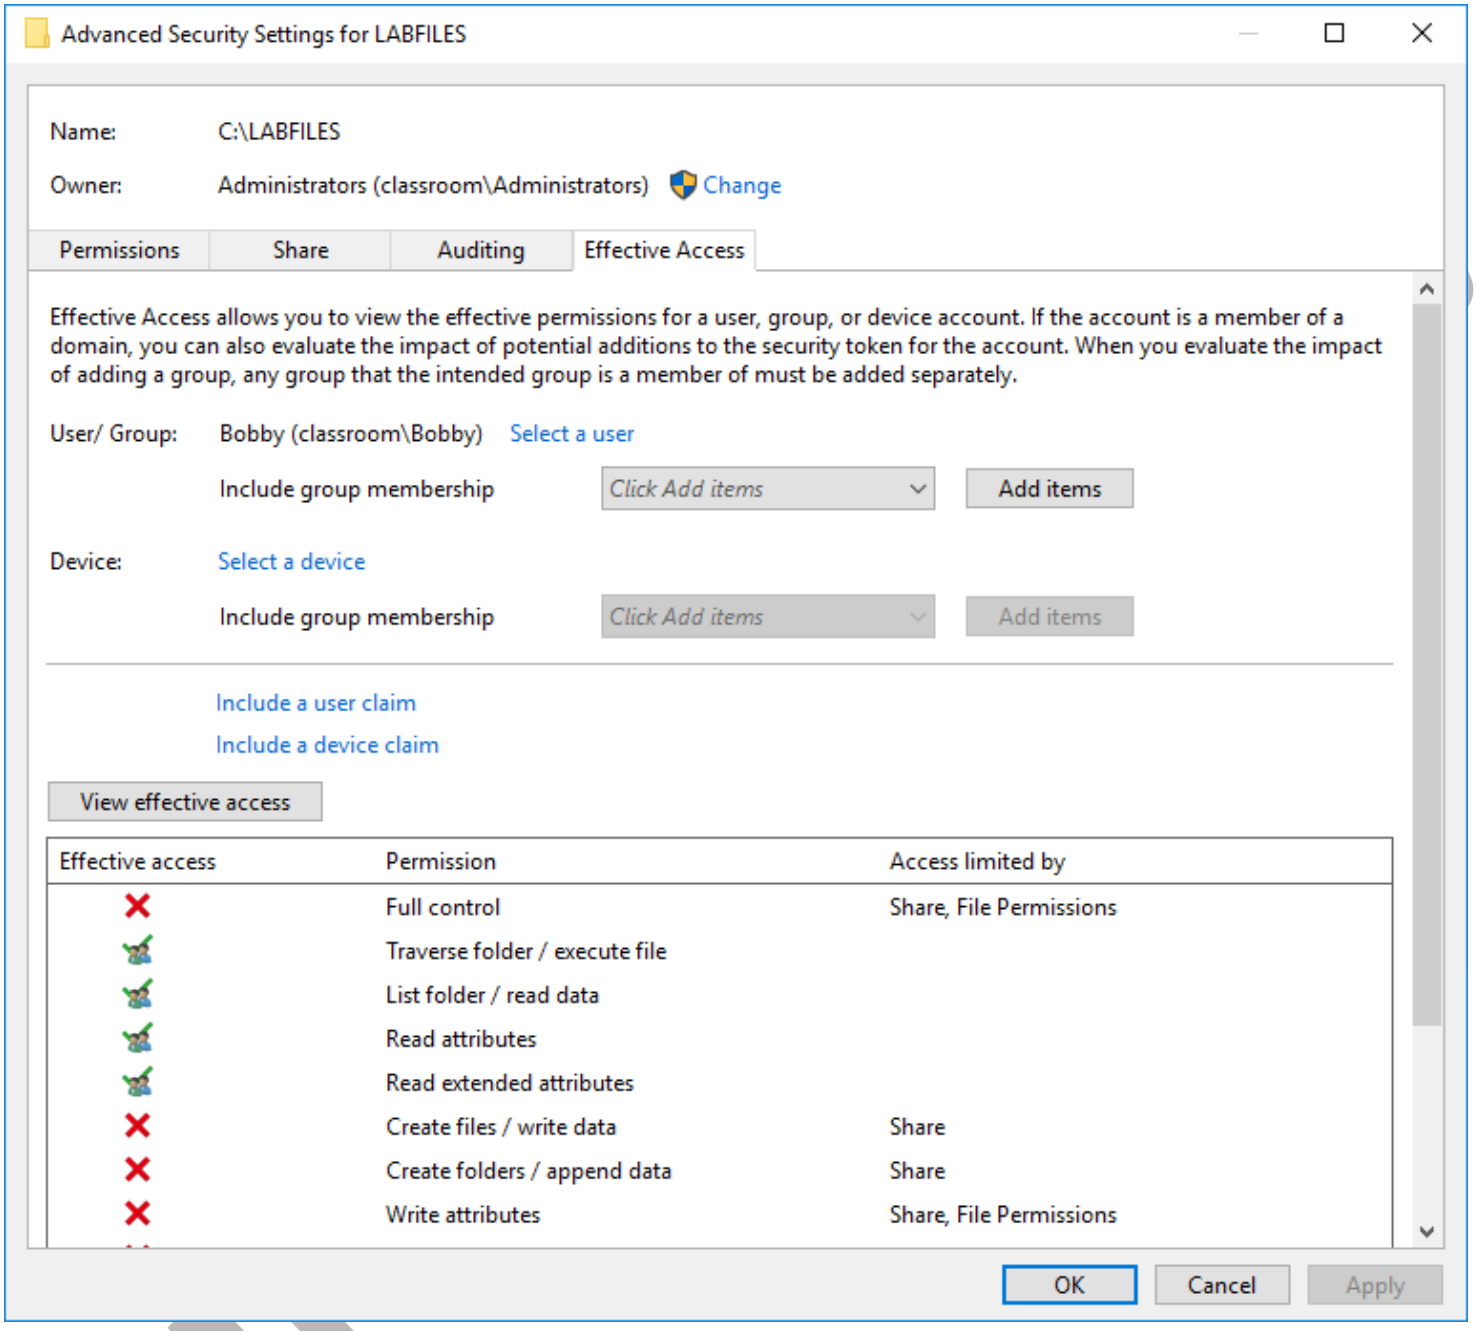

Determining effective permissions for a shared folder. (Screenshot used with permission from Microsoft.)

Provisioning is the process of setting up a service according to a standard procedure or best practice checklist. The IT department must keep track of all assets under management, and user accounts are a type of asset. User accounts are provisioned for new employees and temporary access, such as by consultants and contractors. Some businesses may also need to set up customer accounts.

Provisioning a user account involves the following general steps:

- **Identity proofing** — verifies that the person is who they say they are by checking official documents and records. Circumstances might also demand a background check, which verifies current and previous

addresses, education, or previous employment and whether the person has a criminal record or credit issues.

- **Issuing credentials** — allows the user to select a password known only to them and enroll them with biometric or token-based authenticators.
- **Issuing hardware and software assets** — the user will typically need a computer and smartphone and possibly local copies of licensed software apps. Employees need sufficient resources to do their job. If their resources are inadequate, they might try to obtain hardware and software directly (shadow IT).
- **Teaching policy awareness** — by scheduling training and providing access to learning resources so that the employee or contractor is aware of security policies and risks. They must also be aware of policies for the personal use of any IT assets issued to them.
- **Creating permissions assignment** — by identifying the work roles that the account must support and configuring the appropriate rights using a role-based, mandatory, or attribute-based access control model. If the account is granted privileged access, it should be tagged for close monitoring.

Deprovisioning is the process of removing the access rights and permissions allocated to an employee when they leave the company or from a contractor when a project finishes. This involves removing the account from any roles or security groups. The account might be disabled for a period and then deleted or deleted immediately.

Account Attributes and Access Policies

A user account is defined by a unique security identifier (SID) , a name, and a credential. Each account is associated with a profile. The profile can be defined with custom identity attributes describing the user, such as a full name, email address, contact number, department, etc. The profile may support media such as an account picture.

As well as attributes, the profile will usually provide a location for storing user-generated data files (a home folder). The profile can also store per-account settings for software applications.

Each account can be assigned permissions over files and other network resources and access policies or privileges over the use and configuration of network hosts. These permissions might be assigned directly to the account or inherited through membership in a security group or role. Access policies determine things like the right to log on to a computer locally or via a remote desktop, install software, change the network configuration, etc.

On a Windows Active Directory network, access policies can be configured via group policy objects (GPOs) . GPOs can be used to configure access rights for user/group/role accounts. GPOs can be linked to network administrative boundaries in Active Directory, such as sites, domains, and organizational units (OU).

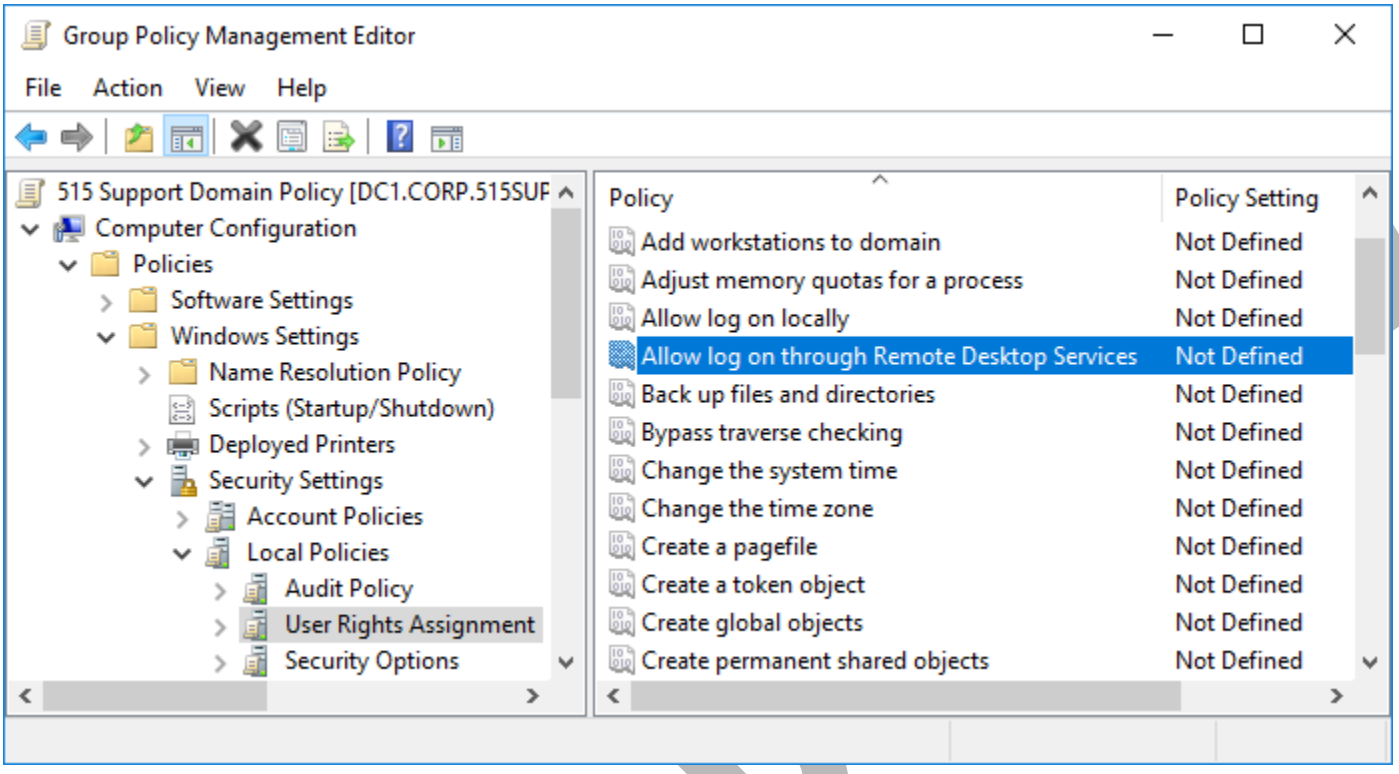

Configuring access policies and rights using Group Policy Objects in Windows Server 2016. (Screenshot used with permission from Microsoft.)

Policy-based restrictions can be used to mitigate some risks of account compromise through the theft of credentials.

Location-Based Policies

A user or device can have a logical network location identified by an IP address, subnet, virtual LAN (VLAN), or organizational unit (OU). This can be used as an account restriction mechanism. For example, a user account may be prevented from logging on locally to servers within a restricted OU.

The geographical location of a user or device can be calculated using a geolocation mechanism:

- **IP address** — can be associated with a map location to varying degrees of accuracy based on information published by the registrant, including name, country, region, and city. The registrant is usually the Internet service provider (ISP), so the information you receive will provide an approximate location of a host based on the ISP. If the ISP is one that serves a large or diverse geographical area, it is more difficult to pinpoint the location of the host Internet service providers (ISPs). Software libraries, such as GeoIP, facilitate querying this data.
- **Location services** — are methods used by the OS to calculate the device's geographical position. A device with a global positioning system (GPS) sensor can report a highly accurate location when outdoors. Location services can also triangulate to cell towers, Wi-Fi hotspots, and Bluetooth signals where GPS is not supported.

Time-Based Restrictions

There are four main types of time-based policies:

- A **time-of-day restrictions policy** establishes authorized login hours for an account.
- A **duration-based login policy** establishes the maximum time an account may be logged in for.
- An **impossible travel time/risky login policy** tracks the location of login events over time. If these do not meet a threshold, the account will be disabled. For example, a user logs in to an account from a device in New York City. A couple of hours later, a login attempt is made from Los Angeles but is refused, and an alert is raised because it is not feasible for the user to be in both locations.
- A **temporary permissions policy** removes an account from a security role or group after a defined period.

4.1.6 Practice Questions (Section Quiz)

q_acc_ctrl_aaa_secp8

An organization's IT department wants to implement a security model responsible for verifying user identities, determining access rights, and monitoring activities within a system.

Which concept is MOST appropriate for the department to implement?

Answers:

- ***AAA**
- Zero trust
- RBAC
- Policy engine

Explanation:

Authentication, authorization, and accounting (AAA) verifies a user's identity (authentication), determines what resources the user can access (authorization), and monitors the user's activities within the system (accounting).

Zero trust might involve elements of AAA, but does not encompass all the aspects of the AAA model.

Role-based access control (RBAC) is the role assigned to individual users within an enterprise. RBAC primarily involves authorization and does not inherently provide all the functionalities of AAA.

The policy engine makeup includes subject and host identities and credentials, access control policies, up-to-date threat intelligence, behavioral analytics, and other results of host and network security scanning and monitoring.

q_acc_ctrl_access_control_01_secp8

The information technology department in a large organization is implementing a new system where the system allows, determines, and enforces various resources based on predefined company guidelines.

Which concept is the department implementing?

Answers:

- ***Policy-driven access control**
- Authorization models
- AAA
- Zero trust

Explanation:

Policy-driven access control uses policies to control access to resources, allowing the organization to systematically enforce rules about who can access which resources under which conditions.

Authorization models are more generic methods for determining what resources a user or system can access within a system. While crucial to any secure system, they do not necessarily incorporate company policies for controlling access.

Authentication, authorization, and accounting (AAA) includes policy enforcement elements in its scope. However, AAA is a broader concept that also involves verifying the user's identity (authentication) and tracking their activities (accounting).

Zero trust does not focus on enforcing company policies to control access, but verifies everything trying to connect to the system.

q_acc_ctrl_access_control_02_secp8

A corporation's IT department is integrating a new framework that permits, ascertains, and applies various resources in accordance with established company policies.

Which principle should the department incorporate?

Answers:

- ***Policy-driven access control**
- Authorization models
- AAA
- Zero trust

Explanation:

Policy-driven access control uses policies to control access to resources, allowing the organization to systematically enforce rules about who can access which resources under which conditions.

Authorization models are more generic methods for determining what resources a user or system can access within a system. While crucial to any secure system, they do not necessarily incorporate company policies for controlling access.

Authentication, authorization, and accounting (AAA) includes policy enforcement elements in its scope. However, AAA is a broader concept that also involves verifying the user's identity (authentication) and tracking their activities (accounting).

Zero trust does not focus on enforcing company policies to control access, but verifies everything trying to connect to the system.

q_acc_ctrl_authentication_secp8

Which of the following is the term for the process of validating a subject's identity?

Answers:

- ***Authentication**
- Identification
- Authorization
- Auditing

Explanation:

Authentication is the process of validating a subject's identity. Authentication includes the identification process, the providing of input by the user to prove his or her identity, and the acceptance of that input as valid by the system.

Authorization grants or denies a subject's access to an object based on the level of permissions or the actions allowed on the object.

Identification identifies the subject. Examples include a username or a user ID number.

Auditing maintains a record of a subject's activity within the information system.

q_acc_ctrl_cia_triad_secp8

What principle of an organization's information security system ensures that only authorized individuals can access sensitive data, the data remains unaltered during storage and transfer, and the data is always accessible when needed?

Answers:

- ***CIA triad**
- Authenticating people
- Access control list
- Two-factor authentication

Explanation:

The CIA (confidentiality, integrity, and availability) triad is critical to information security. Confidentiality ensures that data is accessible only to authorized personnel. Integrity ensures that data remains unaltered during storage and transfer. Availability guarantees that the data is always accessible.

Authenticating people contributes to the confidentiality aspect of the CIA triad by allowing only verified users to access data, but does not cover integrity and availability.

An access control list (ACL) is a rule set that controls network traffic and reduces network attacks. ACL contributes to confidentiality, but does not assure the integrity and availability of data.

Two-factor authentication adds an extra layer of security by requiring a second factor for verification. This does not cover all aspects of the CIA triad.

q_acc_ctrl_gap_analysis_01_secp8

After implementing the National Institute of Standards and Technology (NIST) Cybersecurity Framework, the chief information security officer (CISO) is assessing the company's security posture to identify deficiencies from the framework's recommendations.

What process can the CISO run to get a better sense of what the company needs to improve upon?

Answers:

- Implement business continuity plan
- Penetration test
- Implement disaster recovery plan
- ***Gap analysis**

Explanation:

The CISO would be preparing a gap analysis report. This report shows the defects in the company's current security posture against the NIST Cybersecurity Framework (or any other baseline security framework).

A business continuity plan would significantly improve security posture, but the CISO is verifying what currently exists, not what should exist.

The CISO would not be performing a penetration test. However, this task may show up as a remediation step during a gap analysis report.

Like penetration testing, a disaster recovery plan would likely be a remediation step for any defects the CISO finds during the analysis.

q_acc_ctrl_gap_analysis_02_secp8

The IT department of a corporation evaluates its security mechanisms to identify areas lacking sufficient protection.

Which of the following techniques should the IT department employ?

Answers:

- ***Gap analysis**
- Zero trust
- Authorization models
- Non-repudiation

Explanation:

Gap analysis assesses the differences in performance between a company's information systems or software applications to determine whether they meet requirements. If requirements aren't met, the analysis helps determine what steps to take to meet them.

Zero trust is a strategy that assumes no user or device is trustworthy, whether located inside or outside the network, but does not specifically aim to identify deficiencies in security.

Authorization models dictate what resources a user or system can access within a system. They are not tools used to assess security gaps.

Non-repudiation assures the origin and integrity of transmitted data, preventing entities from denying the validity of the data. However, this is not a technique to identify areas lacking sufficient protection.

q_acc_ctrl_gap_analysis_03_secp8

An IT department is using a technique to assess the differences in performance between their systems, looking to see if the systems meet the established requirements.

Which of the following terms BEST describes the technique the IT department is using?

Answers:

- ***Gap analysis**

- Zero trust
- Authorization models
- Non-repudiation

Explanation:

Gap analysis assesses the differences in performance between a company's information systems or software applications to determine whether they meet requirements. If requirements aren't met, the analysis helps determine what steps to take to meet them.

Zero trust is a strategy that assumes no user or device is trustworthy, whether located inside or outside the network, but does not specifically aim to identify deficiencies in security.

Authorization models dictate what resources a user or system can access within a system. They are not tools used to assess security gaps.

Non-repudiation assures the origin and integrity of transmitted data, preventing entities from denying the validity of the data. However, this is not a technique to identify areas lacking sufficient protection.

q_acc_ctrl_iam_process_secp8

You are the head of the IT department at a large corporation. Recently, there have been several security breaches, and you suspect that these breaches are due to issues with your Identity and Access Management (IAM) processes. You decide to conduct a thorough review of your IAM processes.

Which of the following steps should you prioritize and why?

Answers:

- Reviewing the identification process to ensure that each user, device, or process on the network is uniquely represented.
- Checking the authentication process to verify that each subject is who or what it claims to be when attempting to access a resource.
- ***Assessing the authorization process to determine what rights subjects should have on each resource and whether those rights are being enforced.**
- Examining the accounting process to track authorized usage of a resource or use of rights by a subject and alert when unauthorized use is detected or attempted.

Explanation:

Assessing the authorization process should be the priority. This process determines what rights subjects should have on each resource and whether those rights are being enforced. If this process is flawed, subjects may gain access to resources they should not have access to, leading to security breaches.

Reviewing the identification process is important, but not the first step to take when dealing with security breaches. Identification only ensures that each user, device, or process on the network is uniquely represented, but does not prevent unauthorized access if the authentication and authorization processes are flawed.

Checking the authentication process is crucial, verifying that each subject is who or what it claims to be when attempting to access a resource. However, even if the authentication process is robust, breaches can still occur if the authorization process is flawed and allows subjects to access resources they should not have access to.

Examining the accounting process is important for detecting and alerting when unauthorized use is detected or attempted. However, this process is reactive rather than proactive, helping identify when a breach has occurred, but not preventing breaches from happening in the first place. Therefore, while important, examining the accounting process should not be the first step in addressing security breaches.

q_acc_ctrl_iam_secp8

What is the purpose of identity and access management (IAM) automation in the onboarding process for new employees in an organization?

Answers:

- ***To automate the provisioning and access management tasks associated with new employees.**
- To facilitate knowledge sharing and continuity as employees move into new roles.
- To establish the rules for the acceptable ways in which network and computer systems may be used by defining acceptable behavior by users.
- To carefully plan and assess the implementation of changes in the IT system.

Explanation:

Automating IAM streamlines onboarding, assigns access based on roles and policies, and syncs data with HR systems. This process boosts efficiency, consistency, and security, benefiting organizations in many ways.

Playbooks promote knowledge exchange, consistency, and effectiveness. They guide personnel to maintain operational standards and support transition.

The acceptable use policy (AUP) outlines the appropriate ways to use network and computer systems. It defines acceptable behavior for users.

Assessing and planning changes in an IT system involves a careful evaluation of their impact on related components. However, this process is separate from automating IAM onboarding for new employees.

q_acc_ctrl_nist_framework_secp8

In the context of the NIST Cybersecurity Framework, which function involves identifying, analyzing, containing, and eradicating threats to systems and data security?

Answers:

- Identify
- Protect
- ***Respond**
- Recover

Explanation:

Respond is the correct answer. The Respond function involves taking action regarding a detected cybersecurity incident. The goal is to contain the impact of a potential cybersecurity event.

Identify involves developing an understanding to manage cybersecurity risk to systems, people, assets, data, and capabilities. Identify does not directly involve the process of identifying, analyzing, containing, and eradicating threats.

Protect involves developing and implementing appropriate safeguards to ensure delivery of critical services. Protect is more about prevention than response.

Recover involves maintaining plans for resilience and to restore any capabilities or services that were impaired due to a cybersecurity incident. Recover is more about restoration and recovery after an incident has occurred.

q_acc_ctrl_non_repudiation_01_secp8

When sending confidential data over a network, a company wants to ensure both parties involved cannot deny the validity of the transmitted data.

Which security principle should they prioritize?

Answers:

- ***Non-repudiation**
- Authentication, authorization, and accounting (AAA)
- Adaptive identity
- Zero trust

Explanation:

Non-repudiation is crucial for verifying that transmitted data originated from a verified sender and reached the intended recipient, and neither party can deny the authenticity of the data.

Authentication, authorization, and accounting (AAA) is essential for overall network security but does not provide a specific non-repudiation feature.

Adaptive identity primarily relates to systems that dynamically adjust user access rights based on behavior or other contextual factors. While important for maintaining security, adaptive identity does not directly address non-repudiation.

Zero trust is a security concept centered on the belief that organizations should not automatically trust anything inside or outside their perimeters and instead must verify everything trying to connect to their systems. However, zero trust does not ensure non-repudiation by itself.

q_acc_ctrl_non_repudiation_02_secp8

A company transmits data across a network, ensuring the non-repudiation security principle.

What is the key benefit this provides to both the sender and the recipient of the data?

Answers:

- ***Neither party can deny the authenticity of the data.**
- Both parties can adapt their identity dynamically.
- Both parties should not trust anything inside or outside the network.
- Both parties have control over authentication, authorization, and accounting.

Explanation:

Non-repudiation is a security principle that ensures data originated from a verified sender and reached the intended recipient. This process means neither party can deny the authenticity of the data.

Adaptive identity is a part of advanced access control models but does not directly provide the certainty of data origin and receipt that non-repudiation does.

Not trusting anything inside or outside the network is a zero-trust policy, but does not provide non-repudiation.

While authentication, authorization, and accounting (AAA) is a key security framework used to control user access and track user activity in an organization's network, but does not inherently provide non-repudiation.

q_acc_ctrl_policy_point_secp8

A security analyst wants to ensure that the privileges granted to an individual align with the role within the organization.

What is the primary tool that the analyst should implement?

Answers:

- ***Policy enforcement point**
- Non-repudiation
- Authenticating systems
- Zero trust

Explanation:

Policy enforcement points enforce decisions about whether to grant access to a requested resource or not. Policy enforcement is instrumental in enforcing authorization models.

Non-repudiation ensures that a party to a transaction or communication cannot refute their involvement. However, non-repudiation does not help in aligning individual privileges with roles in an organization.

Authenticating systems does not help specifically with matching privileges to roles, which is the focus of authorization models.

Zero trust is a security concept that recommends not trusting any entity inside or outside the organization by default. Zero trust does not assist in granting access based on roles.

q_acc_ctrl_zero_trust_planes_secp8

You are a cybersecurity expert implementing a zero trust model in a large organization. You are tasked with designing the control and data planes.

Which of the following strategies should you prioritize and why?

Answers:

- Focus on the control plane to ensure that all network devices are properly configured and managed.
- Prioritize the data plane to ensure that data traffic flows securely and efficiently across the network.
- ***Balance your focus between the control and data planes, ensuring both are optimized for security and efficiency.**
- Neither, focus on the application plane to ensure that applications are secure and function properly.

Explanation:

Balancing your focus between the control and data planes is the most effective strategy. In a zero trust model, both planes play crucial roles. The control plane ensures proper network configuration and management, while the data plane handles data traffic. Ensuring both are optimized for security and efficiency can prevent breaches and ensure smooth network operations.

Focusing on the control plane is important as it is responsible for network management and configuration. However, focusing solely on the control plane may leave the data plane, which handles the actual data traffic, vulnerable to attacks. Therefore, while the control plane is crucial, it should not be the only focus.

Prioritizing the data plane is crucial as it handles the actual data traffic across the network. Ensuring the security and efficiency of the data plane can prevent data breaches and ensure smooth network operations. However, without a properly configured control plane, the data plane may not function optimally.

While the application plane is important, it is not the primary focus when designing control and data planes in a zero trust model. The application plane deals with how applications are delivered to end-users and does not directly impact the configuration or security of the control and data planes. Therefore, while it should not be neglected, it should not be the primary focus in this scenario.

q_acc_ctrl_zero_trust_secp8

A network administrator for a technology company is introducing a new cybersecurity model to limit data breaches. They wish to enforce a strategy where every system or user inside or outside the network perimeter must prove their legitimacy before accessing resources.

What principle would be MOST effective in implementing their new strategy?

Answers:

- ***Zero trust**
- Adaptive identity
- Role-based access control
- Policy-driven access control

Explanation:

Zero trust verifies the authenticity of every system or user trying to connect to its resources, serving as the best strategy in this scenario.

Adaptive identity involves systems dynamically adjusting user access rights based on behavior or context rather than requiring every system or user to prove their legitimacy.

Role-based access control (RBAC) does not necessitate the verification of every system or user, which is the central idea of zero trust.

Policy-driven access control does offer a layer of security by determining who to allow access to specific network resources. But policy-driven access control does not incorporate the verify-everything concept of zero trust.

q_acct_bstpract_least_01_secp8

You assign access permissions so that users can only access the resources required to accomplish their specific work tasks.

Which security principle are you complying with?

Answers:

- Job rotation
- ***Principle of least privilege**
- Cross-training
- Need to know

Explanation:

The principle of least privilege is the assignment of access permissions so that users can only access the resources required to accomplish their specific work tasks.

Job rotation and cross-training involve training groups of employees how to perform multiple job roles and periodically rotating those roles. Need to know is a feature of MAC environments where data within your classification level is compartmentalized and requires specific work-task needs for privilege access.

q_acct_bstpract_least_02_secp8

An access control list (ACL) contains a list of users and allowed permissions.

What is it called if the ACL automatically prevents access to anyone who is not on the list?

Answers:

- ***Implicit deny**
- Explicit deny
- Implicit allow
- Explicit allow

Explanation:

With implicit deny, users or groups that are not specifically given access to a resource are denied access. Implicit deny means that there is an assumed or unstated deny that prevents access to anyone not explicitly on the list.

Explicit deny identifies users or objects that are not granted access.

Explicit allow specifically identifies the objects that are allowed access.

Implicit allow is a policy that allows access unless it is explicitly denied. (This ACL type is rarely used.)

q_acct_bstpract_least_03_secp8

You want to implement an access control list in which only the users you specifically authorize have access to the resource. Anyone not on the list should be prevented from having access.

Which of the following methods of access control should the access list use?

Answers:

- ***Explicit allow, implicit deny**
- Explicit allow, explicit deny
- Implicit allow, implicit deny

- Implicit allow, explicit deny

Explanation:

The access list should use explicit allow--users who are allowed access are specifically identified. The access list should also use implicit deny--users who are not explicitly allowed access are denied access.

Explicit deny identifies users or objects that are denied access.

Implicit allow allows access unless it is explicitly denied. This type of ACL is rarely implemented.

q_acct_bstpract_need_secp8

Which of the following principles is implemented in a mandatory access control model to determine object access by classification level?

Answers:

- Separation of duties
- Clearance
- ***Need to know**
- Ownership
- Principle of least privilege

Explanation:

Need to know is used with mandatory access control environments to implement granular control over access to segmented and classified data.

Separation of duties is the security principle that states that no single user is granted sufficient privileges to compromise the security of an entire environment.

Clearance is the subject classification label that grants a user access to a specific security domain in a mandatory access control environment.

Ownership is the access right in a discretionary access control environment that gives a user complete control over an object. This is usually because he or she created the object.

q_acct_bstpract_privilege_secp8

Which of the following is an example of privilege escalation?

Answers:

- ***Privilege creep**
- Principle of least privilege
- Separation of duties
- Mandatory vacations

Explanation:

Privilege creep occurs when a user's job position changes and he or she is granted a new set of access privileges for their new work tasks, but their previous access privileges are not removed. As a result, the user accumulates privileges over time that are not necessary for their current work tasks. This is a form of privilege escalation.

Separation of duties is the security principle that states that no single user is granted sufficient privileges to compromise the security of an entire environment.

Principle of least privilege and separation of duties are countermeasures against privilege escalation.

Mandatory vacations are used to perform peer reviews to help detect mistakes and fraud, and require cross-trained personnel.

q_acct_bstpract_rotation_secp8

You are concerned that the accountant in your organization might have the chance to modify financial information and steal from the company. You want to periodically have another person take over all accounting responsibilities to catch any irregularities.

Which security principle are you implementing by periodically shifting accounting responsibilities?

Answers:

- ***Job rotation**
- Need to know
- Principle of least privilege
- Explicit deny
- Separation of duties

Explanation:

Job rotation is a technique where users are cross-trained in multiple job positions and responsibilities are regularly rotated between personnel. Job rotation can be used for training purposes, but also allows for oversight of past transactions. As jobs rotate, personnel in new positions have the chance to review actions taken by others in that same position and catch security problems.

Separation of duties is the policy of requiring more than one person to complete a task.

The principle of least privilege states that users or groups are given only the access they need to do their job and nothing more.

With explicit deny, users are specifically prevented from gaining access to a resource.

Need to know describes the restriction of data that is highly sensitive and is usually referenced in government and military contexts.

q_acct_bstpract_separate_01_secp8

You want to make sure that any reimbursement checks issued by your company cannot be issued by a single person.

Which security principle should you implement to accomplish this goal?

Answers:

- ***Separation of duties**
- Mandatory vacations
- Job rotation
- Principle of least privilege

Explanation:

Separation of duties is the security principle that states that no single user is granted sufficient privileges to compromise the security of an entire environment.

Job rotation is a technique in which users are cross-trained in multiple job positions and responsibilities are regularly rotated between personnel. Job rotation is used for training purposes, but also allows for oversight of past transactions. As jobs rotate, personnel in new positions have the chance to review actions taken by others in that same position and catch security problems.

A requirement for mandatory vacations requires employees to take vacations of specified length. These vacations can be used to audit actions taken by the employee and provide a passage of time where problems caused by misconduct could become evident.

The principle of least privilege states that users or groups are given only the access they need to do their job and nothing more. With implicit deny, users or groups that are not specifically given access to a resource are denied access.

q_acct_bstpract_separate_02_secp8

Which security principle prevents any one administrator from having sufficient access to compromise the security of the overall IT solution?

Answers:

- Principle of least privilege
- ***Separation of duties**
- Dual administrator accounts
- Need to know

Explanation:

Separation of duties is the security principle that states that no single user is granted sufficient privileges to compromise the security of an entire environment. Usually, this principle is implemented by dividing administrative privileges among several administrators.

The principle of least privilege states that users should have the minimal amount of access necessary to perform their work tasks.

A dual-administrator accounts policy ensures that each administrator has a privileged-level account and a normal user-level account.

Need to know is an access control tool used in mandatory access control environments to implement granular control over access to segmented and classified data.

q_acct_bstpract_separate_03_secp8

What is the primary purpose of separation of duties?

Answers:

- Increase the difficulty of performing administrative duties.
- ***Prevent conflicts of interest.**
- Inform managers that they are not trusted.
- Grant a greater range of control to senior management.

Explanation:

The primary purpose of separation of duties is to prevent conflicts of interest by dividing administrative powers between several trusted administrators. This prevents a single person from having all of the privileges over an environment, which would create a primary target for attack and a single point of failure.

Increasing administrative difficulty, informing managers that they are not trusted, or granting a greater range of control to senior management are not the primary purposes of separation of duties. Separation of duties might seem to increase administrative difficulty, but this separation provides significant security benefits.

A manager is informed they are not trusted when they are not given any responsibility as opposed to a reasonable portion of responsibility.

Senior management already has full control over their organization.

q_acct_bstpract_something_you_have_secp8

After a breach, an organization implements new multi-factor authentication (MFA) protocols.

What MFA philosophy incorporates using a smart card or key fob to support authentication?

Answers:

- ***Something you have**
- Something you are
- Somewhere you are
- Something you know

Explanation:

Something you have means the account holder possesses something that no one else does, such as a smart card, key fob, or smartphone that can generate or receive a cryptographic token.

Something you are refers to a biometric or inherence factor. A biometric factor uses physiological identifiers, such as a fingerprint or facial scan, or behavioral identifiers, such as how someone moves (gait).

Somewhere you are means the system applies a location-based factor to an authentication decision. Location-based authentication measures some statistics about where you are.

Something you know means the information used for authentication is from something one can recall, such as a passphrase or username/password combination.

q_acct_bstpract_something_you_know_secp8

A cyber engineer conducts a multi-factor authentication (MFA) assessment of an organization's authentication security.

What MFA philosophy uses knowledge factors and includes passphrases to gain access to systems?

Answers:

- ***Something you know**
- Something you have
- Something you are
- Somewhere you are

Explanation:

Something you know means the information used for authentication is from something one can recall, such as a passphrase or username/password combination.

Something you have means the account holder possesses something that no one else does, such as a smart card, key fob, or smartphone that can generate or receive a cryptographic token.

Something you are refers to a biometric or inherence factor. A biometric factor uses physiological identifiers, such as a fingerprint or facial scan, or behavioral identifiers, such as how someone moves (gait).

Somewhere you are means the system applies a location-based factor to an authentication decision. Location-based authentication measures some statistics about where you are.

q_acc_models_abac_01_secp8

Which access control model is based on assigning attributes to objects and using Boolean logic to grant access based on the attributes of the subject?

Answers:

- ***Attribute-based access control (ABAC)**
- Mandatory access control (MAC)
- Role-based access control (RBAC)
- Rule-based access control

Explanation:

The ABAC model is based on assigning attributes to objects and using Boolean logic to grant access based on the attributes of the subject.

The MAC model is based on classification labels being assigned to objects and clearance labels being assigned to subjects. When a subject's clearance lines up with an objects classification, the subject is granted access.

The RBAC model grants access based on the subject's role in an organization.

The Rule-Based Access Control model grants access based on a set of rules or policies.

q_acc_models_abac_02_secp8

The IT department of an international organization is responsible for managing access controls to various sensitive resources and systems. The IT department aims to select the access control model that aligns BEST with the organization's security requirements, user roles, and data sensitivity levels.

Which access control model enforces permissions based on attributes and predefined rules, allowing the IT department to make fine-grained access decisions based on users' characteristics and data sensitivity within the multinational corporation?

Answers:

- ***Attribute-based access control (ABAC)**
- Rule-based access controls
- Discretionary access control (DAC)
- Role-based access control (RBAC)

Explanation:

Attribute-based access control (ABAC) enforces access permissions based on attributes, such as user attributes, resource attributes, and environmental conditions, allowing for fine-grained access control decisions.

Rule-based access controls rely on predefined rules without considering attributes, leading to less granular access control than ABAC.

Discretionary access control (DAC) gives users more control over resource access, which may not be suitable for managing fine-grained access decisions based on data sensitivity and user characteristics.

Role-based access control (RBAC) focuses on associating permissions with roles, not attributes, which is less suitable for fine-grained access control based on users' characteristics and data sensitivity within the organization.

q_acc_models_abac_03_secp8

A large multinational company uses a cloud-based document storage system. The system provides access to documents by considering a combination of factors: the user's department, geographic location, the document's sensitivity level, and the current date and time.

For example, only the finance department of a specific region can access its financial reports, and they can do so only during business hours.

Which access control model does the company MOST likely use to manage this complex access control?

Answers:

- ***Attribute-based access control (ABAC)**
- Rule-based access controls
- Discretionary access control (DAC)
- Role-based access control (RBAC)

Explanation:

Attribute-based access control (ABAC) enforces access permissions based on attributes, such as user attributes, resource attributes, and environmental conditions, allowing for fine-grained access control decisions.

Rule-based access controls (RBAC) rely on predefined rules without considering attributes, leading to less granular access control than attribute-based access control (ABAC).

Discretionary access control (DAC) gives users more control over resource access, which may not be suitable for managing fine-grained access decisions based on data sensitivity and user characteristics.

Role-based access control (RBAC) focuses on associating permissions with roles, not attributes, which is less suitable for fine-grained access control based on users' characteristics and data sensitivity within the organization.

q_acc_models_dac_01_secp8

You have a system that allows the owner of a file to identify users and their permissions to the file.

Which type of access control model is implemented?

Answers:

- ***Discretionary access control (DAC)**
- Mandatory access control (MAC)
- Role-based access control (RBAC)
- Rule-based access control

Explanation:

This is an example of a discretionary access control list (DACL), which uses the discretionary access control (DAC) model. With DAC, individuals use their own discretion (decisions or preferences) for assigning permissions and allowing or denying access.

Mandatory access control (MAC) uses labels for both subjects (users who need access) and objects (resources with controlled access). When a subject's clearance lines up with an object's classification, and when the user has a need to know (referred to as a category), the user is granted access.

Role-based access control (RBAC) allows access based on a role in an organization, not individual users. Roles are defined based on job description or a security access level. Users are made members of a role and receive the permissions assigned to the role.

Rule-based access control uses characteristics of objects or subjects along with rules to restrict access. Access control entries identify a set of characteristics that are examined for a match. If all characteristics match, access is either allowed or denied based on the rule.

q_acc_models_dac_02_secp8

Which form of access control enforces security based on user identities and allows individual users to define access controls over owned resources?

Answers:

- Role-based access control (RBAC)
- Mandatory access control (MAC)
- ***Discretionary access control (DAC)**
- Attribute-based access control (ABAC)

Explanation:

Discretionary access control (DAC) uses identities to control resource access. Users can make their own decisions about how much access to grant to other users.

Role-based access control (RBAC) allows access based on a role in an organization.

Mandatory access control (MAC) uses labels for both subjects and objects.

Attribute-based access control (ABAC) restricts access by assigning attributes to resources.

q_acc_models_least_privilege_01_secp8

In a medium-sized tech company, employees have different roles and responsibilities requiring access to specific resources and data. The IT team is implementing security measures to control access effectively and reduce the risk of unauthorized activities.

What security measure could the IT team implement in the tech company to control access effectively and minimize the risk of unauthorized activities?

Answers:

- ***The principle of least privilege, granting each employee the minimum necessary access based on job roles.**
- Implement a firewall to protect the company's network from external threats.
- Enforce mandatory password changes every month to enhance password security.
- Implement intrusion detection systems to monitor and identify potential security breaches.

Explanation:

The IT team is implementing the principle of least privilege, granting each employee the minimum necessary access based on job roles to control access effectively and minimize the risk of unauthorized activities.

Although implementing a firewall is an essential security measure, the scenario does not mention it as part of the access control measures.

Enforcing mandatory password changes can improve password security, but the scenario does not mention it as one of the measures implemented to control access.

Implementing intrusion detection systems is crucial for monitoring and identifying potential security breaches, but the scenario does not specify it as part of the access control measures.

q_acc_models_least_privilege_02_secp8

A new hire has just joined the IT department of a large organization. The human resources (HR) department assigns the employee an initial set of credentials for accessing the company network.

However, the new hire requires additional access to other systems within the company network.

Given this context, which of the following represents the MOST accurate example of implementing the principle of least privilege in this organization?

Answers:

- ***The IT department provides access to only a specific system requested by the new hire.**
- The IT department automatically grants the new hire access to all systems upon joining.
- The IT department grants full admin rights to the new hire after their first request.
- The IT department waits for a request from HR before granting any access to the new hire.

Explanation:

The principle of least privilege states that users should only have access to the resources necessary for their job roles.

Automatically granting the new hire access to all systems upon joining is against the principle of least privilege and leads to unnecessary security risks.

Granting full admin rights to the new hire after the first request is against the principle of least privilege. New hires should only have access to systems necessary for their jobs.

Waiting for a request from HR before granting any access does not necessarily follow the principle of least privilege. The IT department should base the decision on the requirements of the new hire's job role, not on a request from HR.

q_acc_models_least_privilege_03_secp8

What is the purpose of implementing the principle of least privilege in endpoint protection?

Answers:

- To restrict access to specific network resources.
- To enforce mandatory security configurations on devices.
- To manage firewall rules across an organization's network.
- ***To grant users, applications, and processes only the minimum necessary permissions.**

Explanation:

The purpose of implementing the principle of least privilege (PoLP) in endpoint protection is to grant users, applications, and processes the minimum necessary permissions required to perform their specific duties and tasks.

Restricting access to specific network resources is a general access control measure, but it is not specifically related to the principle of least privilege.

Enforcing mandatory security configurations on devices and managing firewall rules across an organization's network are both security measures for network security but do not relate directly to the principle of least privilege.

q_acc_models_least_privilege_04_secp8

One of your company's accountants submitted a ticket stating they could not access a particular section of the accounting software.

Why might the accountant not have access to every part of the accounting software?

Answers:

- Licensing
- DAC
- MAC
- ***Least privilege**

Explanation:

To increase the security posture of any given system, users should only have the necessary access (least privilege) to complete their work and nothing more.

Depending on the software, not having a specific license prevents users from accessing all of a given software. However, the most common reason is least privilege.

The resource owner has primacy in a discretionary access control (DAC) model. Every resource has an owner who creates a file or service, although another user can receive ownership assignment.

Security clearance levels form the basis of mandatory access control (MAC). Rather than defining access control lists (ACLs) on resources, each object receives a classification label. Depending on the clearance level, a subject receives access to that resource.

q_acc_models_least_privilege_05_secp8

A security analyst at a large organization aims to minimize the attack surface. To reach this goal, the analyst seeks to reduce the vulnerabilities an attacker can exploit, decrease the amount of code in use, and limit system interactions.

Which strategy should the security analyst implement to achieve this objective effectively?

Answers:

- Install the latest antivirus software
- Increase the complexity of user passwords
- ***Implement the principle of least privilege**
- Replace the wireless network with a wired network

Explanation:

Implementing the principle of least privilege ensures that users, systems, and processes have the minimum privileges needed to perform their tasks, thus reducing the attack surface.

Installing the latest antivirus software helps to detect and remove malicious software but does not necessarily limit the number of vulnerabilities an attacker can exploit.

Increasing the complexity of user passwords is part of a robust password policy, but it does not reduce the attack surface.

Although replacing the wireless network with a wired network might reduce certain attack vectors related to wireless security, it does not address the broader concept of minimizing the number of vulnerabilities, amount of code, or system interactions.

q_acc_models_least_privilege_06_secp8

A company acquires a smaller company and has its in-house technical team review the new systems before allowing them on the existing network.

During this review, the technical team discovers users with unnecessary permissions, user accounts for former employees, and no longer needed groups.

These discoveries indicate the violation of what BEST practice?

Answers:

- ***Principle of least privilege**
- Access control
- File system permissions
- Configuration enforcement

Explanation:

The principle of least privilege dictates that users, applications, and processes have the minimum permissions necessary to complete their duties and nothing more. This best practice helps ensure the security of the entire network.

Access control relates to the permissions granted to individuals, software, networks, or systems to allow one of these entries to complete authorized actions. It is not a best practice.

File system permissions are the permissions granted to each file based on an access control list (ACL). It is not a best practice.

Configuration enforcement ensures that systems and devices within a network have the same configuration, allowing for easier identification of compromised systems and ensuring a standard baseline.

q_acc_models_least_privilege_07_secp8

A financial services company tasks its IT security team with reducing the network's attack surface. They have segmented the network into security zones, put port security measures in place, and physically isolated critical servers.

The IT security team wants to further reduce the risk of attack by managing traffic flow between security zones.

Which of the following measures should the team implement?

Answers:

- ***Apply the principle of least privilege when defining traffic policies between zones.**
- Implement MAC filtering on all switch ports.
- Establish an air-gapped network for all company servers.
- Use a star network topology.

Explanation:

Least privilege can effectively reduce the attack surface by only allowing necessary communications between security zones, minimizing potential points of exploitation.

Media access control (MAC) filtering can help secure individual switch ports, but it does not directly affect the attack surface between security zones.

Air-gapped networks provide isolation but are impractical for most servers in companies that need to communicate with other networks. Also, it does not address traffic control between different security zones.

A flat or star network topology does not reduce the attack surface between security zones, but could increase attack surface within zones by allowing unrestricted exchange between nodes.

q_acc_models_location-based_secp8

An employee traveling in Europe for vacation submitted a ticket as they could not access their work email.

Which policy does the company use?

Answers:

- Password management
- Password age
- Multi-factor authentication
- ***Location-based authentication**

Explanation:

Location-based access policies would need a temporary exemption option to allow for travel. Location-based access policies prevent access to company systems outside a specified area (typically the company's state).

Password management would not aid in accessing a system set up to prevent access outside specific locations. However, it would provide a better way for employees to save passwords securely.

Password age would not help resolve issues accessing corporate resources from outside specific approved locations.

Multi-factor authentication would not help resolve issues with accessing corporate resources from outside specific approved locations.

q_acc_models_mac_01_secp8

Which type of access control focuses on assigning privileges based on security clearance and data sensitivity?

Answers:

- ***Mandatory access control (MAC)**
- Role-based access control (RBAC)
- Task-based access control (TBAC)
- Discretionary access control (DAC)

Explanation:

Mandatory access control (MAC) uses classifications to assign privileges based on security clearances and data sensitivity.

Role-based access control (RBAC) is a form of access control that assigns privileges based on a job description. New users are simply assigned a job label. The job label holds all the privileges necessary to accomplish the work tasks assigned to that job.

Task-based access control (TBAC) defines individual work tasks to assign privileges. It is similar to RBAC but with a primary difference. The difference is that a single user may be assigned dozens of tasks under TBAC, while under RBAC each user is only assigned a single job description.

With Discretionary access control (DAC), an administrator or owner defines user and resource access.

q_acc_models_mac_02_secp8

In which form of access control environment is access controlled by rules rather than identity?

Answers:

- Discretionary access control (DAC)
- ***Mandatory access control (MAC)**
- Access control list (ACL)
- Most client-server environments

Explanation:

A mandatory access control (MAC) environment controls access based on rules rather than identity.

Discretionary access control (DAC) environments use identity to control access.

Access control lists (ACLs) are a specific example of an identity-based access control mechanism used in DAC environments.

Most client-server environments use ACLs and, therefore, are DAC solutions.

q_acc_models_mac_03_secp8

The IT department of a governmental agency manages access controls for its various systems and resources. It is currently evaluating different access control models to enhance security across the enterprise. The organization deals with sensitive data, and it is crucial to have proper controls in place to protect the organization from unauthorized access.

The IT team considered several access control models, each offering distinct features, and the IT team analyzed which one aligns BEST with the organization's security requirements.

Which access control model enforces access permissions based on data sensitivity and predefined security labels, ensuring a higher level of control over sensitive resources in the organization?

Answers:

- ***Mandatory access control (MAC)**
- Discretionary access control (DAC)
- Role-based access control (RBAC)
- Attribute-based access control (ABAC)

Explanation:

Mandatory access control (MAC) enforces access permissions based on data sensitivity and predefined security labels. MAC restricts users' ability to change access permissions, ensuring higher resource control.

Discretionary access control (DAC) allows users to determine access permissions to resources owned. It does not focus on data sensitivity or predefined security labels.

Role-based access control (RBAC) assigns access permissions based on job roles and responsibilities. While it streamlines access management, it does not enforce control based on data sensitivity or predefined security labels.

Attribute-based access control (ABAC) enforces access permissions based on user characteristics, resource properties, and environmental factors. It provides a flexible and dynamic access control mechanism but does not explicitly focus on data sensitivity.

q_acc_models_mac_04_secp8

An international defense organization has developed a classified software system used for satellite communication. The system processes various levels of classified information: top secret, secret, and confidential.

The access to this system and its data is not determined by the individual user's wishes, a role, or by evaluating multiple attributes and conditions. Instead, the system enforces system access based on predefined classifications attached to subjects (users) and objects (data).

Which access control model is the defense organization MOST likely using?

Answers:

- ***Mandatory access control (MAC)**
- Discretionary access control (DAC)
- Role-based access control (RBAC)
- Attribute-based access control (ABAC)

Explanation:

Mandatory access control (MAC) enforces access permissions based on data sensitivity and predefined security labels. MAC restricts users' ability to change access permissions, ensuring higher resource control.

Discretionary access control (DAC) allows users to determine access permissions to resources owned. DAC does not focus on data sensitivity or predefined security labels.

Role-based access control (RBAC) assigns access permissions based on job roles and responsibilities. While RBAC streamlines access management, it does not enforce control based on data sensitivity or predefined security labels.

Attribute-based access control (ABAC) enforces access permissions based on user characteristics, resource properties, and environmental factors. ABAC provides a flexible and dynamic access control mechanism but does not explicitly focus on data sensitivity.

q_acc_models_provisioning_secp8

A recently hired information technology manager wants to implement more automation regarding the onboarding procedure.

What process describes setting up accounts so a new employee can automatically access the software and file shares from the human resource platform?

Answers:

- Multi-factor authentication
- Following least privilege
- Enabling a password reuse policy
- ***Provisioning**

Explanation:

Provisioning is the process of setting up a service according to a standard procedure or best practice checklist. Linking multiple systems together can increase the automation of onboarding procedures.

Multi-factor authentication is an authentication scheme that requires the user to present at least two different factors as credentials.

To increase the security posture of any given system, users should only have the necessary access (least privilege) to complete their work and nothing more.

The National Institute of Standards and Technology (NIST) recommends that people refrain from using passwords across multiple systems or services. However, this would not increase automation.

q_acc_models_rbac_01_secp8

Which form of access control is based on job descriptions?

Answers:

- Discretionary access control (DAC)
- Mandatory access control (MAC)
- ***Role-based access control (RBAC)**
- Attribute-based access control (ABAC)

Explanation:

RBAC is based on job descriptions.

DAC is based on identity.

MAC is based on rules.

ABAC restricts access by assigning attributes to resources.

q_acc_models_rbac_02_secp8

You have implemented an access control method that only allows users who are managers to access specific data.

Which type of access control model is being used?

Answers:

- ***Role-based access control (RBAC)**
- Mandatory access control (MAC)
- Discretionary access control (DAC)
- Discretionary access control list (DACL)

Explanation:

Role-based access control (RBAC) allows access based on a role in an organization, not individual users. Roles are defined based on job description or a security-access level. Users are made members of a role and receive the permissions assigned to the role.

Discretionary access control (DAC) assigns access directly to subjects based on the discretion of the owner. Objects have a discretionary access control list (DACL) with entries for each subject. Owners add subjects to the DACL and assign rights or permissions. The permissions identify the actions the subject can perform on the object.

Mandatory access control (MAC) uses labels for both subjects (users who need access) and objects (resources with controlled access). When a subject's clearance lines up with an object's classification, and when the user has a need to know (referred to as a category), the user is granted access.

q_acc_models_rule_based_secp8

In a medium-sized company, the IT department manages access to various systems and resources for employees. The team wants to enhance the security posture by implementing better access controls. They use rule-based access controls and time-of-day restrictions to achieve this goal.

What are the IT department's objectives in implementing rule-based access controls and time-of-day restrictions? (Select two.)

Answers:

- ***To define specific access rules based on employees' roles and responsibilities.**
- ***To restrict access to critical systems during non-working hours to enhance security.**
- To ensure all employees have access to all resources at any time for increased productivity.
- To eliminate the need for user authentication and simplify access management.
- To ensure that all employees are biometrically verified.

Explanation:

The IT department aims to implement rule-based access controls to define specific access rules based on employees' roles and responsibilities, ensuring users can access only the resources they need to fulfill their job duties and enhance security.

The implementation of time-of-day restrictions limits access to critical systems during non-working hours, which helps improve security by reducing exposure to potential threats when fewer employees are present.

Allowing all employees to access all resources without restrictions poses significant security risks, violates the principle of least privilege, and increases the attack surface.

Implementing rule-based access controls and time-of-day restrictions does not aim to eliminate user authentication; rather, it complements user authentication mechanisms.

Biometrics are a type of authentication and do not apply to rule-based access controls.

q_acc_models_rule_secp8

Which of the following is an example of rule-based access control?

Answers:

- ***Router access control lists that allow or deny traffic based on the characteristics of an IP packet.**
- A subject with a government clearance that allows access to government classification labels of Confidential, Secret, and Top Secret.
- A member of the accounting team that is given access to the accounting department documents.
- A computer file owner who grants access to the file by adding other users to an access control list.

Explanation:

A router access control list that allows or denies traffic based on the characteristics of an IP packet is an example of rule-based access control.

A subject with a government clearance that allows access to government classification labels of Confidential, Secret, and Top Secret is an example of mandatory access control.

A member of the accounting team that is given access to the accounting department documents is an example of role-based access control.

A computer file owner who grants access to the file by adding other users to an access control list is an example of discretionary access control.

q_acc_models_teach_awareness_secp8

You are a system administrator for a large corporation. A new employee, John, has just joined the marketing team and you are tasked with provisioning his user account.

John is not very familiar with the company's IT policies and procedures.

What should be your first step in this process?

Answers:

- Assign John the same privileges as the rest of the marketing team.
- ***Teach John about the company's IT policies and the importance of policy awareness.**
- Create a user account for John with full administrative privileges.
- Provide John with a list of all the software he will need to install on his computer.

Explanation:

Before provisioning John's account, it's crucial to ensure he understands the company's IT policies and the importance of policy awareness. This will help him understand his responsibilities, the limitations of his account, and the potential consequences of policy violations. This is the correct answer as it emphasizes the importance of policy awareness before account provisioning.

Assigning John the same privileges as the rest of the marketing team might seem like a logical step, but without understanding the company's IT policies, John might misuse or not fully understand his privileges, which could lead to security risks.

Creating a user account for John with full administrative privileges is not a recommended practice. This could lead to serious security risks, especially if John is not familiar with the company's IT policies.

Providing John with a list of all the software he will need to install on his computer is a task that should be done after his account has been provisioned and he has been made aware of the company's IT policies. This is not the first step in the process.

q_acc_models_time-based_01_secp8

A contractor only works for a company from 9 a.m. to 12 p.m.

What kind of restriction could the company set up on the contractor's account to prevent using it outside that range?

Answers:

- Location-based restrictions
- Password restrictions
- ***Time-based restrictions**
- Mandatory access control

Explanation:

A time-based restriction would prevent access to corporate resources outside a set schedule. The company should set this account policy.

Location-based access policies would need a temporary exemption option to allow for travel. Location-based access policies prevent access to company systems outside a specified area (typically the company's state).

Password policies would not help prevent access to systems for the contractor during the slated schedule.

Security clearance levels form the basis of mandatory access control (MAC). Rather than defining access control lists (ACLs) on resources, each object receives a classification label. Depending on the clearance level, a subject receives access to that resource.

q_acc_models_time-based_02_secp8

As a network security administrator, you have implemented various access control models in your organization.

One day, you notice that an employee's account has been accessed from two different geographical locations within a time frame that would be impossible for the employee to travel.

The locations are New York and Tokyo, and the logins occurred within an hour of each other.

What should be your immediate course of action?

Answers:

- Ignore the logins as it might be a system error.
- ***Immediately disable the account and investigate the logins.**
- Contact the employee and ask if they shared their login credentials.
- Change the password of the employee's account.

Explanation:

Disabling the account immediately prevents any further potential unauthorized access. Investigating the logins can help determine if it was a case of credential sharing, a system error, or a potential security breach. This is the correct course of action as it addresses the issue of impossible travel time/risky login policy.

Ignoring the logins could potentially allow unauthorized access to continue, which could lead to data breaches or other security incidents. This is not a recommended course of action.

While contacting the employee is a good step, it should not be the immediate action. The account should first be secured to prevent any potential unauthorized access. Furthermore, if the employee has indeed shared their credentials, they might not be truthful about it.

Changing the password might prevent further unauthorized access, but it does not address the potential security breach that has already occurred. Also, if the unauthorized user has access to the employee's email, they could potentially reset the password again.

4.2 Authentication

As you study this section, answer the following questions:

- What is the difference between authentication factors and attributes?
- What is an example of the "something you are" authentication type?
- What is an example of the "something you have" authentication type?
- What is multi-factor authentication?
- Which physical attributes can be used to identify an individual?

In this section, you will learn to:

- Use a biometric scanner
- Use single sign-on

The key terms for this section include:

Term	Definition
Multi-factor authentication (MFA)	An authentication scheme that requires the user to present at least two different factors as credentials; for example, something you know, something you have, something you are, something you do, and somewhere you are. Specifying two factors is known as 2FA.
Factors	In authentication design, different technologies for implementing authentication, such as knowledge, ownership/token, and biometric/inherence. These are characterized as something you know/have/are.
Personal identification number (PIN)	A number used in conjunction with authentication devices such as smart cards; as the PIN should be known only to the user, loss of the smart card should not represent a security risk.
Hard authentication token	Authentication token generated by a cryptoprocessor on a dedicated hardware device. As the token is never transmitted directly, this implements an ownership factor within a multi-factor authentication scheme.
Smart cards	A security device similar to a credit card that can store authentication information, such as a user's private key, on an embedded cryptoprocessor.
One-time password (OTP)	A password that is generated for use in one specific session and becomes invalid after the session ends.

Term	Definition
Security key	Portable HSM with a computer interface, such as USB or NFC, used for multi-factor authentication.
Soft authentication token	OTP sent to a registered number or email account or generated by an authenticator app as a means of two-step verification when authenticating account access.
Passwordless	Multi- factor authentication scheme that uses ownership and biometric factors, but not knowledge factors.
Attestation	Capability of an authenticator or other cryptographic module to prove that it is a root of trust and can provide reliable reporting to prove that a device or computer is a trustworthy platform.
NT LAN Manager (NTLM) authentication	A challenge-response authentication protocol created by Microsoft for use in its products.
Pluggable authentication module (PAM)	A framework for implementing authentication providers in Linux.
Directory service	A network service that stores identity information about all the objects in a particular network, including users, groups, servers, client computers, and printers.
Lightweight Directory Access Protocol (LDAP)	Protocol used to access network directory databases, which store information about authorized users and their privileges, as well as other organizational information.
Distinguished name (DN)	A collection of attributes that define a unique identifier for any given resource within an X.500-like directory.
Single sign-on (SSO)	Authentication technology that enables a user to authenticate once and receive authorizations for multiple services.
Kerberos	A single sign-on authentication and authorization service that is based on a time-sensitive, ticket-granting system.
Key distribution center (KDC)	A component of Kerberos that authenticates users and issues tickets (tokens).

Term	Definition
Ticket Granting Ticket (TGT)	In Kerberos, a token issued to an authenticated account to allow access to authorized application servers.
Federation	A process that provides a shared login capability across multiple systems and enterprises. It essentially connects the identity management services of multiple systems.
Identity provider (IdP)	In a federated network, the service that holds the user account and performs authentication.
Security Assertion Markup Language (SAML)	An XML-based data format used to exchange authentication information between a client and a service.
Simple Object Access Protocol (SOAP)	An XML-based web services protocol that is used to exchange messages.
Representational State Transfer (REST)	A standardized, stateless architectural style used by web applications for communication and integration.
Open Authorization (OAuth)	A standard for federated identity management, allowing resource servers or consumer sites to work with user accounts created and managed on a separate identity provider.
JavaScript Object Notation (JSON)	A file format that uses attribute-value pairs to define configurations in a structure that is easy for both humans and machines to read and consume.
Biometric authentication	An authentication mechanism that allows a user to perform a biometric scan to operate an entry or access system. Physical characteristics stored as a digital data template can be used to authenticate a user. Typical features used include facial pattern, iris, retina, fingerprint pattern, and signature recognition.
False Rejection Rate (FRR)	A biometric assessment metric that measures the number of valid subjects who are denied access.
False Acceptance Rate (FAR)	A biometric assessment metric that measures the number of unauthorized users who are mistakenly allowed access.
Crossover Error Rate (CER)	A biometric evaluation factor expressing the point at which FAR and FRR meet, with a low value indicating better performance.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts.
	• Confidentiality, Integrity, and Availability (CIA)
	• Non-repudiation
	• Authentication, Authorization, and Accounting (AAA)
	  ○ Authenticating people
	  ○ Authenticating systems
	  ○ Authorization models
	• Gap analysis
	• Zero trust
	  ○ Control plane
	  ▪ Adaptive identity
	  ▪ Threat scope reduction
	  ▪ Policy-driven access control
	  ▪ Policy Administrator
	  ▪ Policy Engine
	  ○ Data plane
	  ▪ Implicit trust zones
	  ▪ Subject/System
	  ▪ Policy enforcement point
	
	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Access control
	• Least privilege
	
	4.6 Given a scenario, implement and maintain identity and access management.
	• Provisioning/de-provisioning user accounts
	• Permission assignments and implications
	• Identity proofing
	• Access controls
	  ○ Mandatory
	  ○ Discretionary
	  ○ Role-based
	  ○ Rule-based
	  ○ Attribute-based
	  ○ Time-of-day restrictions
	  ○ Least privilege
	• Multi-factor authentication
	  ○ Implementations
	  ▪ Hard/soft authentication tokens
	  ○ Factors
	  ▪ Something you know
	  ▪ Something you have
	  ▪ Something you are
	  ▪ Somewhere you are

Exam	Objective
	5.1 Summarize elements of effective security governance. • Policies ○ Information security policies

4.2.1 Authentication (Lesson Video)

Transcript:

In this video, we'll discuss authentication. Authentication involves verifying a claimed identity. To verify an identity, we need some unique information or data that could only come from that individual.

Three main factors are used in authentication: something you know, something you have, and something you are.
First, let's look at something you know. This commonly used authentication method is knowledge-based. It requires the user to provide information that only they know. The most familiar example is a password. The password is something that only the user should know for their authentication.
A PIN is another example of something you know authentication. Unlike a password, a PIN is limited to numerical digits.
Another example is cognitive information. This is when a question is provided that only the user can answer. Often, for account recovery methods, a company will ask users to supply answers to specific questions, such as your mother's maiden name or the name of your first pet.
The next authentication type, something you have, is sometimes called token-based authentication. Hard tokens require the user to possess a physical factor—an object such as a smart card, key fob, or a USB device—that can be used for authentication. Soft tokens are software that can be installed on a user's smartphone, tablet, or computer.
Soft tokens generate one-time passwords (OTPs) that are used for authentication. When users attempt to log in, they must enter their username and password and then generate or receive an OTP from the soft token application. The OTP is time-limited and changes periodically, usually every 30 seconds. This adds an extra layer of security since OTP is unique for each login attempt and cannot be easily replicated or reused by an attacker.
The third factor, something you are, refers to physiological attribute biometrics. Your fingerprint, retinal scan, vein scan, face scan, or voice print will prove that you are who you say you are.
On their own, some of these authentication methods could be easily compromised. A password can be written down, guessed, or shared. A smart card can be lost or stolen. When two or more of these factors are combined, this is known as multifactor authentication. When the requirement for a password is combined with a token-generated PIN or combined with a fingerprint scan, accessing an account becomes much more complicated. With Multifactor authentication, a username and password can be breached but are unusable without the additional factor!
While we've talked about the three main authentication factors, other factors can be used to help validate authentication: somewhere you are, something you do, something you exhibit, and someone you know. While these may not verify a user's identity independently, they help improve security and work well in multifactor authentication settings.
Somewhere you are authentication uses a physical location to verify a user's identity. For example, you may have a desktop system configured to allow authentication requests only if the user has passed through the building's entrance using their ID card. If the user tries to log in to the desktop, the system will look to see if the user has entered the building. If not, the user's account is locked.
The location could be determined using radio-frequency identification or an RFID proximity reader. If the user is within range, they're authenticated to the system. The user moving out of range automatically locks the system.
GPS location data can also be used to unlock mobile devices. For example, your phone can be automatically unlocked within a specific geographical area. If it moves outside that location, it will be locked and require a PIN or password to unlock. Bluetooth can also be used for this type of authentication. If you're in range of a device via Bluetooth, you can configure your mobile device to unlock itself automatically.
Another authentication type is something you do. This type requires the user to perform a particular action to verify the user's identity. For example, the user might be required to supply a handwriting sample analyzed against a baseline

sample. It could also require a typing test, where the user must type some sample text, and their typing behaviors are analyzed against a baseline sample.
Another attribute could be something you exhibit, such as a personality trait or a habit. Authentication systems can note the time of day you usually log on, your access method, or the tasks you often perform. When administrators notice unusual activity or risky behavior, they may restrict access to these users. This could mean requiring a password change, another authentication method, or blocking user access.
Another authentication factor is someone you know. You've probably heard the phrase, "It's not what you know, it's who you know." Having someone who can vouch for you can go a long way to establishing relationships and building trust. The same is true for authentication. An example of this trust is certificates.
Certificate authorities issue certificates to organizations that they trust. These organizations are intermediary certificate authorities who then issue certificates to their users. This creates a chain of trust that verifies that the subject named in the certificate can be trusted and has been verified by a certificate authority. Certificates are commonly used for logging into secure sites or signing documents electronically.
That's it for this lesson. In this lesson, we discussed the three main factors of authentication. These include something you know, something you have, and something you are. We also discussed other authentication factors, such as somewhere you are, something you do, something you exhibit, and someone you know.

4.2.2 Authentication Factors Facts

Assuming that an account has been created securely and the identity of the account holder has been verified, authentication verifies that only the account holder is able to use the account and that the system may only be used by account holders. Authentication technologies allow the use not only of passwords but also of biometric and token factors to better secure accounts. Understanding the strengths and weaknesses of these factors will help you to implement and maintain strong authentication systems.

This lesson covers the following topics:

- Authentication design
- Multifactor authentication
- Authentication factors

Authentication Design

To access resources on a network, you must prove who you are and have the required permissions. This process consists of the following elements:

- *Identification* is the initial process of confirming your identity when you request credentials. It occurs when you enter a user ID to log on. Identity proofing occurs during the identification phase as you prove that you are who you say you are to obtain credentials. Suppose you have been identified previously but cannot provide the assigned authentication credentials (such as a lost password). In that case, identity proofing is called upon again.
- *Authentication* is the verification of the issued identification credentials. It is usually the second step in the identification process. It establishes your identity, ensuring that you are who you say you are.

Identity is as simple as telling someone your name. In the computer world, a username is a form of identification. Identification alone is not very secure because anyone could pretend to be you. To substantiate your identity, you need to provide some verification that you are who you say you are. The following chart provides a few of the basics of identity authentication.

Term	Description
Identity provider (IdP)	An identity provider is an online service that manages identity information for other organizations. The IdP creates records from an organization's existing data and policies. These records are used to authenticate user requests.
Attributes	Attributes can be your role, position, or current project. This information can be used to determine policy and permission.
Certificates	Certificates are issued by a certificate authority and verify identity by providing the following: • Public keys • Details on the owner of the certificate • Details on the issuer of the certificate
Tokens	A token is a device or a file used to authenticate. A hardware token, such as a key fob, serves as something you have. A software token (or a soft token) is stored in devices such as laptops, desktops, or mobile phones. These tokens are specific to the device and cannot be altered or duplicated.
SSH keys	A secure shell (SSH) key is an access credential. It operates like usernames and passwords but is mainly used to implement single sign-on and other automated processes.

Authentication is performed when a supplicant or claimant presents credentials to an authentication server. The server compares what was presented to the copy of the credentials it has stored. If they match, the account is authenticated. *Authentication design* refers to selecting a technology that meets requirements for confidentiality, integrity, and availability:

- **Confidentiality** , in terms of authentication, is critical because if account credentials are leaked, threat actors can impersonate the account holder and act on the system with whatever rights they have.
- **Integrity** means that the authentication mechanism is reliable and not easy for threat actors to bypass or trick with counterfeit credentials.
- **Availability** means that the time taken to authenticate does not impede workflows and is easy enough for users to operate.

Multifactor Authentication

An authentication design that uses only passwords or a single knowledge factor is considered weak. Password secrets are too prone to compromise to be reliable. Other types of authentication factors can be used to supplement or replace password-based logins. A multifactor authentication (MFA) technology combines the use of more than one type of factor.

Multifactor authentication requires a combination of different technologies. For example, requiring a PIN along with a date of birth may be stronger than entering a PIN alone, but it is not multifactor.

You might also see references to two-factor authentication (2FA) . This just means that there are precisely two factors involved, such as an ownership-based smart card or biometric identifier with something you know, such as a password or PIN.

Authentication Factors

There are many different technologies for defining credentials. These can be categorized as factors .

Something You Know Factor

The longest-standing authentication factor is "Something You Know," or a knowledge factor. The typical knowledge factor is the *login* , composed of a username and a password. This is the weakest type of authentication because both items are something you know, but this is also the most used.

The username is typically not a secret (although it should not be published openly), but the password must be known only to the account holder. Only the passwords or other information associated with the usernames can be used to validate identity. A passphrase is a longer password composed of several words. This has the advantages of being more secure and easier to remember. Composition passwords are created by the system and are usually two or more unrelated words divided by symbols on the keyboard. A final example is cognitive information, such as questions you can only answer, such as your mother's maiden name, the model of your first car, or the city where you were born.

Windows sign-in screen. (Screenshot used with permission from Microsoft.)

A personal identification number (PIN) is also something you know. Originally, PINs were associated with short four- or six-digit numeric sequences used with bank cards. In modern authentication designs, the main characteristic of a PIN is that it is valid for authenticating to a single device only. This type of PIN can use any character and be any length.

Something You Are Factor

Something you are refers to a biometric or inherence factor. A *biometric factor* uses either physiological identifiers, such as a fingerprint or facial scan, or behavioral identifiers, such as the way someone moves (gait). The identifiers are scanned and recorded as a template. When the user authenticates, another scan is taken and compared to the template. This is one of the more secure forms of authentication and one of the most expensive.

Somewhere You Are Factor

Somewhere you are (also known as geolocation) means the system applies a location-based factor to an authentication decision. Location-based authentication measures some statistics about where you are. This could be a physical geographic location measured using a device's location service or Internet Protocol (IP) network address. A device's IP address could be used to refer to a logical network segment, or it could be linked to a geographic location using a geolocation service. Within a premises network, the physical port location, virtual LAN (VLAN), or Wi-Fi network can also be made the basis of location-based authentication.

Location-based authentication is not used as a primary authentication factor, but it may be used as a continuous authentication mechanism or as an access control feature. For example, if a user enters the correct credentials at a remote access gateway, but their IP address shows them to be in a different country than expected, access controls might be applied to restrict the privileges granted or refuse access completely. Another example is when a user appears to log-in from multiple geographic locations that would be physically impossible with travel time.

Examples of implementations include:

- A desktop system configured to allow authentication requests only if you have passed through the building's entrance using your ID card. If you are not in the building, your account is locked.
- A system configured with an RFID proximity reader and required RFID badges. Authentication requests are allowed if you are within the workstation's RFID range. If you move out of range, the workstation is immediately locked. Reauthentication is not allowed until you move back within range.
- GPS location data is used to determine a device's location. Authentication requests are allowed if you and the device are in a specified location. If not, the device is locked, or additional authentication factors are requested.
- Wi-Fi triangulation is used to determine a device's location. Authentication requests are allowed if you and the device are in a specified location. If not, the device is locked, or additional authentication factors are requested.

Something You Can Do Factor

Something you can do requires performing a particular action to verify your identity. Here are a few examples of an action that can be used:

- Supply a handwritten sample that’s analyzed against a baseline sample for authentication.
- Type sample text. Your typing behaviors are analyzed against a baseline before authentication.

Something You Exhibit Factor

Something that you exhibit could include a personality trait or a habit. For example:

- The time of day you usually log on.
- The method you typically use to access information.
- The types of tasks you usually perform.

When administrators notice unusual or risky behavior, they may restrict access. This could mean requiring a password change, another authentication method, or blocking access.

Someone You Know Factor

Having someone who can vouch for you can go a long way in establishing relationships and building trust. The same is valid with authentication. Certificates and attestation are examples of this attribute.

Something You Have Factor

Something you have is an *ownership factor* . It means that the account holder possesses something that no one else does. The term token-based authentication is often used to describe this factor due to the use of security tokens. Some examples include:

- Swipe cards (like credit cards) with authentication information stored on the magnetic strip.
- Photo IDs are handy when combined with other forms of authentication. Still, they are high risk if they are the only form of required authentication. Photo IDs are easily manipulated or reproduced, require personnel for verification, and cannot be verified against a system.
- Key fobs are small, programmable hardware often used to access buildings and open doors. Key fobs are usually attached to a keychain.
- Security tokens generate a unique password when activated manually. These passwords are used once and usually expire in minutes.
- Smart cards contain a memory chip with encrypted authentication information.
- Smartphone that can generate or receive a cryptographic token.

An ownership factor means that the user possesses some type of device that only they can operate. This is referred to as an authenticator. The authenticator is able to generate or receive a token that identifies and authenticates the user. There are three main types of token generation:

- **Certificate-based authentication** — is when the supplicant controls a private key that can generate a unique signed token. The identity provider can verify the signature via the public key. The main drawback of this approach is the administrative burden of implementing PKI to issue digital certificates.
- **One-time password (OTP)** — is when a token is generated using some sort of hash function on a shared secret value plus a synchronization seed, such as a timestamp (TOTP) or HMAC (HOTP). The token can only be used once. A new token is generated for each authentication decision. This approach still uses a key pair and hashing for security, but it does not require PKI.
- **Fast IDentity Online (FIDO) Universal 2nd Factor (U2F)** — uses a public/private key pair to register each account, avoiding the need to communicate a shared secret, a weakness of HOTP and TOTP. The private key is locked to the U2F device and signs the token; the public key is registered with the authentication server and verifies the token. As no digital certificates are involved, the solution does not rely on PKI.

A hard authentication token is generated within a secure cryptoprocessor. The authentication design means that there is no transmission of the token itself. Several device-based authenticators can be used to implement hard tokens:

- **Smart cards** — implement certificate-based authentication. A certificate is a digital document associated with a user as a one-to-one or many-to-one mapping. In a one-to-one mapping, each certificate maps to an individual user account (each user has a unique certificate). With many-to-one mapping, a certificate maps to many user accounts (a group of users shares the same certificate). The smart card stores the user's digital certificate, the private key associated with the certificate, and a personal identification number (PIN) used to activate the card. The card must be presented to a reader. There are physical contact and contactless near-field communication (NFC) card types.
- **One-time password (OTP)** — refers to a cryptoprocessor that can generate a token. This type of hardware token does not need an interface to connect with a computer; the user just reads the code displayed.

- **Security key** — refers to a portable hardware security module (HSM) with a computer interface, such as USB or NFC. They are most closely associated with U2F, but some might also support certificate-based authentication or HOTP/TOTP. A security key must be activated to show presence. Some keys just have an activation button, but most use a biometric fingerprint reader for better security. A PIN must also be configured as a backup mechanism.

Key fob token generator. (Image © 123RF.com.)

There are also simpler smart cards and fobs that simply transmit a static token programmed into the device. For example, many building entry systems work on the basis of static codes. These mechanisms are highly vulnerable to cloning and replay attacks.

A soft authentication token is a one-time password generated by the identity provider and transmitted to the supplicant. The OTP could be sent to a registered phone number as an SMS/text message or to an email account. This method is more likely to use counter-based tokens, though they will still have an expiry period.

Soft tokens sent via SMS or email do not really count as an ownership factor. These systems can be described as two-step verification rather than MFA. The tokens are highly vulnerable to interception.

A more secure soft OTP token can be generated using an authenticator app. This is software installed on a computer or smartphone. The user must register each identity provider with the app, typically using a scannable quick response (QR) code to communicate the shared secret. When prompted to authenticate, the user must unlock the authenticator app with their device credential to view the OTP token. There is less risk of interception than with an SMS or email message, but as it runs on a shared-use device, there is the possibility that malware could compromise the app.

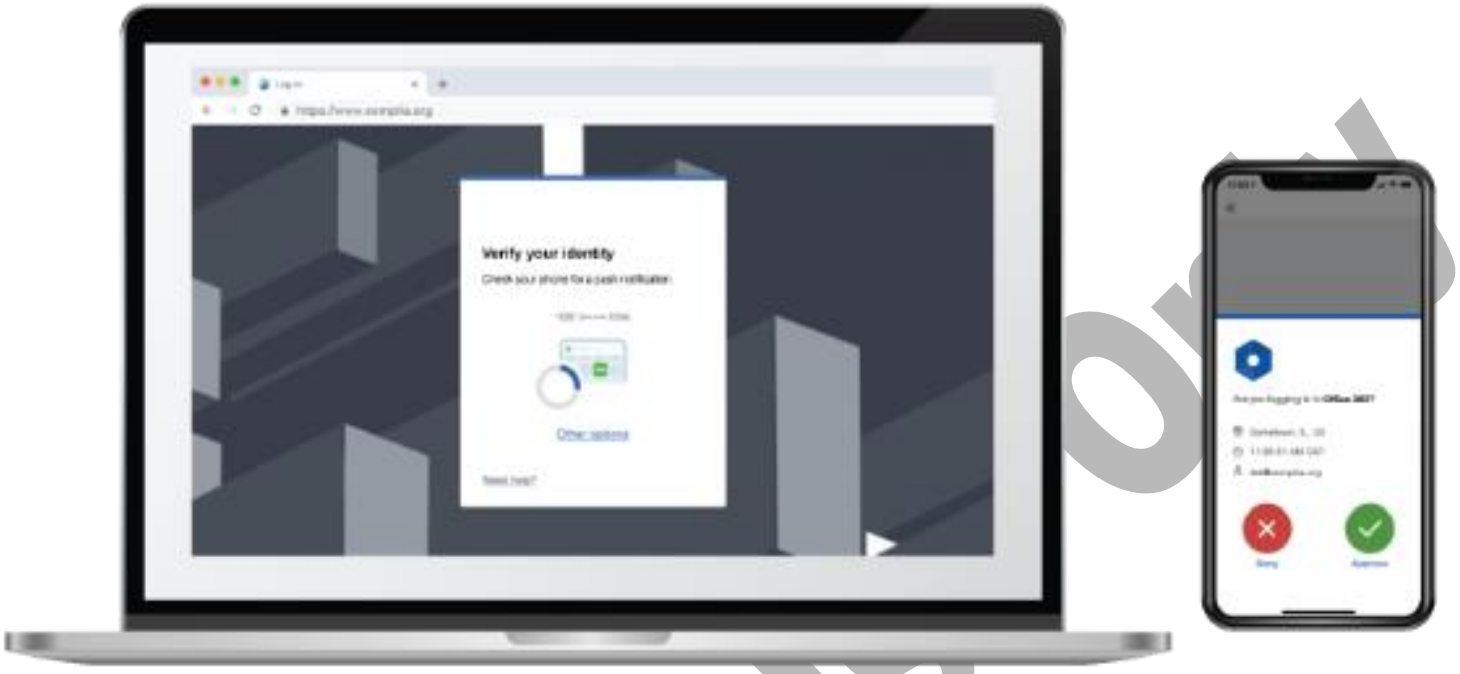

Using an authenticator app to sign in to a site. After the user signs in with a password, the site prompts them to authorize the sign-in using the authenticator app installed on their smartphone.

With token-based MFA, the user account is typically still configured with a password. This might be used as a backup mechanism, or there might be a two-step verification process where the user must enter their password and then submit an OTP.

Passwordless means that the whole authentication system no longer processes knowledge-based factors. The FIDO2 with WebAuthn specifications provides a framework for passwordless authentication. It works as follows:

- The user chooses either a roaming authenticator, such as a security key, or a platform authenticator implemented by the device OS, such as Windows Hello or Face ID/Touch ID for macOS and iOS.
- The user configures a secure method or local gesture to confirm presence and authenticates the device. This gesture could be a fingerprint, face recognition, or PIN. This credential is only ever validated locally by the authenticator.
- The user registers with a web application or service, referred to as a relying party. For each new relying party, the authenticator generates a public/private key pair. The user's client browser obtains the public key from the authenticator and registers it to associate it with an account on the relying party.
- When presented with an authentication challenge, the user performs the local gesture to unlock the private key. The private key is used to sign a confirmation that the local gesture worked, which is then sent to the relying party.
- The relying party uses the public key to verify the signature and authenticate the account session.

As with FIDO U2F, this provides similar security to smart card authentication but does not require accounts to have digital certificates and PKI, reducing the management burden. FIDO2 WebAuthn improves on FIDO U2F by adding an application programming interface (API) that allows web applications to work without a password element for authentication. Most FIDO U2F authenticators should also support FIDO2/WebAuthn.

For a passwordless system to be secure, the authenticator must be trusted and resistant to spoofing or cloning attacks. **Attestation** is a mechanism for an authenticator device, such as a FIDO security key or the TPM in a PC or laptop, to prove that it is a root of trust. Each security key is manufactured with an attestation and model ID. During the registration step, if the

relying party requires attestation, the authenticator uses this key to send a report. The relying party can check the attestation report to verify that the authenticator is a known brand and model and supports whatever cryptographic properties the relying party demands.

Note that the attestation key is not unique; if it were unique, it would be easy to identify individuals and be a serious threat to privacy. Instead, it identifies a particular brand and model.

4.2.3 Authentication Methods (Lesson Video)

Transcript:

A user must provide the correct credentials to authenticate a system, usually a username and password. After that, the system checks an access service or protocol to ensure a match with existing credentials. The authentication process itself doesn't determine your authorization level. In other words, authentication doesn't determine what you can access; it only defines who you are. The application determines your authorization level and what you can do with the application only after you authenticate to the system. Let's look at a few authentication methods you need to know as a security professional.

Single sign-on, or SSO, is an authentication process that allows users to access multiple systems, applications, websites, and other resources using only a single set of credentials. With SSO, a user authenticates once using designated credentials and can access different resources seamlessly.
Imagine that you need to log into your workstation, log into the company portal, log in to check your emails, and then log in again to access your video conferencing software. Instead of assigning four different credentials, it would be better to implement an SSO. Since this service shares authentication sessions between systems, you automatically receive access to all the sessions if you log into one.
While SSO dramatically improves usability, it comes with the risk that breached credentials can be used to gain access to a wide array of resources. Considering SSO provides powerful, seamless access to a wide range of sensitive systems and data using only a single set of credentials, multifactor authentication methods should be coupled with SSO to prevent credentials from being easily abused and stolen.
Directory services provide a single sign-on location for multiple network resources. Examples include Microsoft's Active Directory and LDAP directory services. Directory services' users sign on with a domain user account to gain access to available domain resources.
Another authentication method is OAuth. Although OAuth works closely with authentication services, it doesn't provide authentication services itself. Instead, OAuth provides authorization services. In other words, it authorizes an entity to obtain information but doesn't verify the entity itself. OAuth lets a third party access a user's information from another website without a password. It acts as an intermediary between two entities and provides access tokens that manage the exchange of specific information.
OAuth has a few main components. It starts with a user wanting to access multiple apps without creating many new accounts. Next, you need the client software app on a computer, mobile phone, or any other smart device. You also need a resource server that the user and client are trying to access and an authorization server that provides a token for the user and client to access the resource server.
Let's say that I just installed an app onto my phone. When I open the new app, I'm prompted to log in using my Facebook or Google account. I decide to use my Facebook account, so I'm redirected to Facebook's authorization server to get an access token. The authorization server authenticates you as the user, and the app uses OAuth to grant an access token with precise information from the app. Now, when the app goes to the resource server with that token, you're registered and logged into the app with limited access to your Facebook information.
So, where's the authentication? Well, that's where OpenID Connect, or OIDC, comes in. OIDC is built on the OAuth framework and provides the needed authentication service. Remember, OAuth only provides authorization. OIDC takes it a step further and provides identity authentication. OIDC also accesses stored personal information and preferences from an identity layer that the user has previously disclosed. This layer provides the website with the user's information, where, when, and how of the authentication, and which attributes the user wants to share with that specific website and why. Big companies like Google and Microsoft use OIDC, so it's secure and reliable.
Another authentication method is a federation. A federation is a group of domains that established enough trust to share authorizations with each other. A federation can be used within a single organization with multiple domains or include

several trusted organizations that want to pool their resources. The good thing about this authentication method is that everything's on-site, providing the administration with detailed access control levels.
Another authentication method is attestation. Attestation is a protocol that proves that software can be trusted. It tells the remote user that the software is legitimate and certified. Attestation usually works in both parties' interest.
For example, say I was going to log into my bank account. I'd want to be sure that the site that I'm logging into is trustworthy, and the bank would like to be sure that I am really the one who is trying to log into my account.
That's it for this lesson. In this lesson, we talked about several different authentication methods. These include single sign-on, directory services, Oauth, federation, and attestation. As a security professional, knowing and understanding these authentication methods will benefit you immensely in your future endeavors.

4.2.4 Authentication Methods Facts

While an on-premises network can use a local directory to manage accounts and rights, as organizations move services to the cloud, these authorizations have to be implemented using federated identity management solutions.

This lesson covers the following topics:

- Local, network, and remote authentication
- Windows authentication
- Linux authentication
- Directory services
- Single sign-on authentication
- Federation
- Open authorization

Local, Network, and Remote Authentication

One of the most important features of an operating system is the *authentication provider* , which is the software architecture and code that underpins the mechanism by which the user is authenticated before starting a shell.

Knowledge-based authentication relies on cryptographic hashes. A plaintext password is not usually transmitted or stored in a credential database because of the risk of compromise. Instead, the password is stored as a cryptographic hash. When a user enters a password to log in, an authenticator converts what is typed into a hash and transmits that to an authority. The authority compares the submitted hash to the one in the database and authenticates the subject only if they match.

Windows Authentication

Windows authentication involves a complex architecture of components (docs.microsoft.com/en-us/windows-server/security/windows-authentication/credentials-processes-in-windows-authentication), but the following three scenarios are typical:

- **Windows local sign-in** — is the Local Security Authority Subsystem Service (LSASS) that compares the submitted credential to a hash stored in the Security Accounts Manager (SAM) database, which is part of the registry. This is also referred to as *interactive logon* .
- **Windows network sign-in** — is LSASS, which can pass the credentials for authentication to an Active Directory (AD) domain controller. The preferred system for network authentication is based on Kerberos, but legacy network applications might use NT LAN Manager (NTLM) authentication .
- **Remote sign-in** — is used if the user's device is not directly connected to the local network. Authentication can take place over a virtual private network (VPN), enterprise Wi-Fi, or web portal. These use protocols to create a secure connection between the client machine, the remote access device, and the authentication server.

Linux Authentication

In Linux, local user account names are stored in /etc/passwd . When a user logs in to a local interactive shell, the password is checked against a hash stored in /etc/shadow . Interactive login over a network is typically accomplished using Secure Shell (SSH). With SSH, the user can be authenticated using cryptographic keys instead of a password.

A pluggable authentication module (PAM) is a package for enabling different authentication providers, such as smart-card log-in. The PAM framework can also be used to implement authentication to network directory services.

Directory Services

A directory service stores information about users, computers, security groups/roles, and services. Each object in the directory has a number of attributes. The directory schema describes the types of attributes, what information they contain, and whether they are required or optional. In order for products from different vendors to be interoperable, most directory services are based on the Lightweight Directory Access Protocol (LDAP) , which was developed from a standard called X.500.

Within an X.500-like directory, a distinguished name (DN) is a collection of attributes that define a unique identifier for any given resource. A distinguished name is made up of attribute-value pairs separated by commas. The most specific attribute is listed first, and successive attributes become progressively broader. This most specific attribute is the relative distinguished name, as it uniquely identifies the object within the context of successive (parent) attribute values.

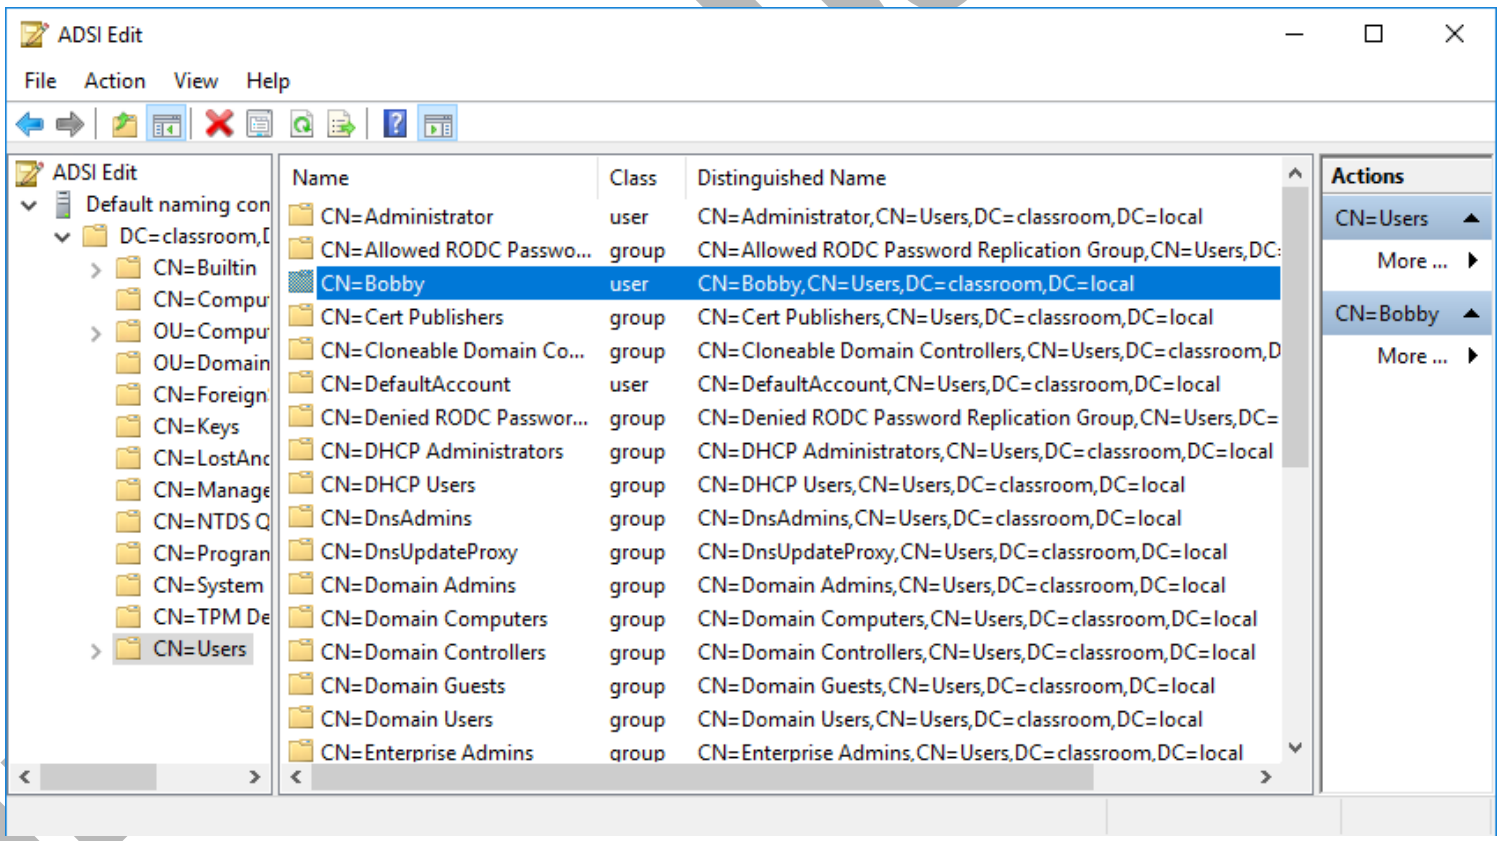

Browsing objects in an Active Directory LDAP schema. (Screenshot used with permission from Microsoft.)

Some of the attributes commonly used include common name (CN), organizational unit (OU), organization (O), country (C), and domain component (DC).

For example, the distinguished name of a web server operated by Widget in the UK might be the following:

CN=WIDGETWEB, OU=Marketing, O=Widget, C=UK, DC=widget, DC=foo

Single Sign-on Authentication

A single sign-on (SSO) system allows the user to authenticate once and then receive authorizations on compatible application servers without having to enter credentials again.

Kerberos is a single sign-on network authentication and authorization protocol used on many networks, notably as implemented by Microsoft's Active Directory (AD) service. Kerberos was named after the three-headed guard dog of Hades (Cerberus) because it consists of three parts. Clients request services from application servers, which rely on an intermediary—a key distribution center (KDC) —to vouch for their identity. There are two services that make up a KDC: the Authentication Service and the Ticket Granting Service.

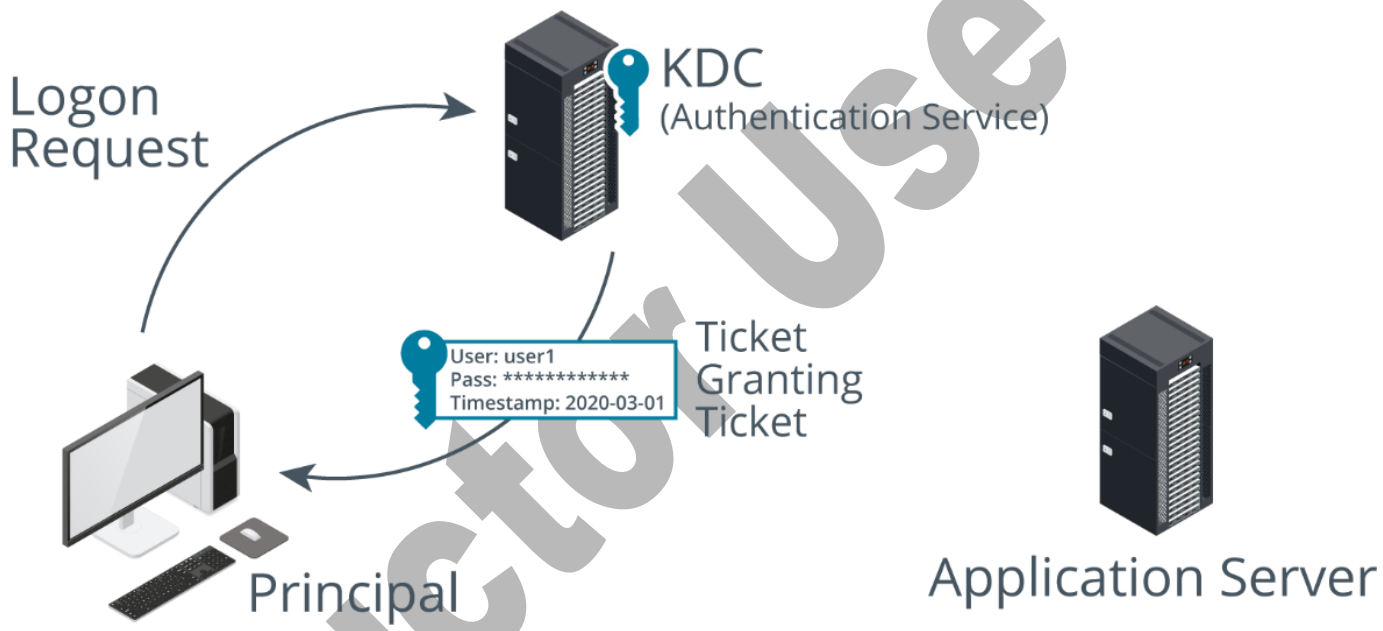

Kerberos Authentication Service. (Images © 123RF.com.)

Kerberos can authenticate human users and application services. These are collectively referred to as *principals* . Using authentication to a Windows domain as an example, the first step in Kerberos SSO is to authenticate with a KDC server implemented as a domain controller.

- The principal sends the authentication service (AS) a request for a Ticket Granting Ticket (TGT) . This is composed by encrypting the date and time on the local computer with the user's password hash as the key.

The password hash itself is not transmitted over the network. Although we refer to passwords for simplicity, the system can use other authenticators, such as smart card login. The AS checks that the user account is present, that it can decode the request by matching the user's password hash with the one in the Active Directory database, and that the request has not expired. If the request is valid, the AS responds with the following data:

- **Ticket Granting Ticket (TGT)** — contains information about the client (name and IP address) plus a time stamp and validity period. This is encrypted using the KDC's secret key.

- **TGS session key** — communicates between the client and the Ticket Granting Service (TGS). This is encrypted using a hash of the user's password.

The TGT is an example of a logical token. All the TGT does is identify who you are and confirm that you have been authenticated—it does not provide you with access to any domain resources.

Federation

Federation is the notion that a network needs to be accessible to more than just a well-defined group of employees. In business, a company might need to make parts of its network open to partners, suppliers, and customers. The company can manage its employee accounts easily enough. Managing accounts for each supplier or customer internally may be more difficult. Federation means the company trusts accounts created and managed by a different network. As another example, in the consumer world, a user might want to use both Google Workspace and Twitter. If Google and Twitter establish a federated network for the purpose of authentication and authorization, then the user can log on to Twitter using their Google credentials or vice versa.

An on-premises network can use technologies such as LDAP and Kerberos, very often implemented as a Windows Active Directory network, because the administration of accounts and devices can be centralized. When implementing federation, authentication and authorization design comes with more constraints and additional requirements to ensure interoperability between different platforms. Web applications might not support Kerberos, while third-party networks might not support direct federation with Active Directory/LDAP. The design for these cloud networks likely requires the use of other standard protocols or frameworks for interoperability between web applications.

These interoperable federation protocols use claims-based identity. While the technical implementation and terminology are different, the overall model is similar to that of Kerberos SSO:

- The principal attempts to access a service provider (SP). The service provider redirects the principal to an identity provider (IdP) to authenticate.
- The principal authenticates with the identity provider and obtains a claim in the form of some sort of token or document signed by the IdP.
- The principal presents the claim to the service provider. The SP can validate that the IdP has signed the claim because of its trust relationship with the IdP.
- The service provider can now connect the authenticated principal to its own accounts database to determine its permissions and other attributes. It may be able to query attributes of the user account profile held by the IdP if the principal has authorized this type of access.

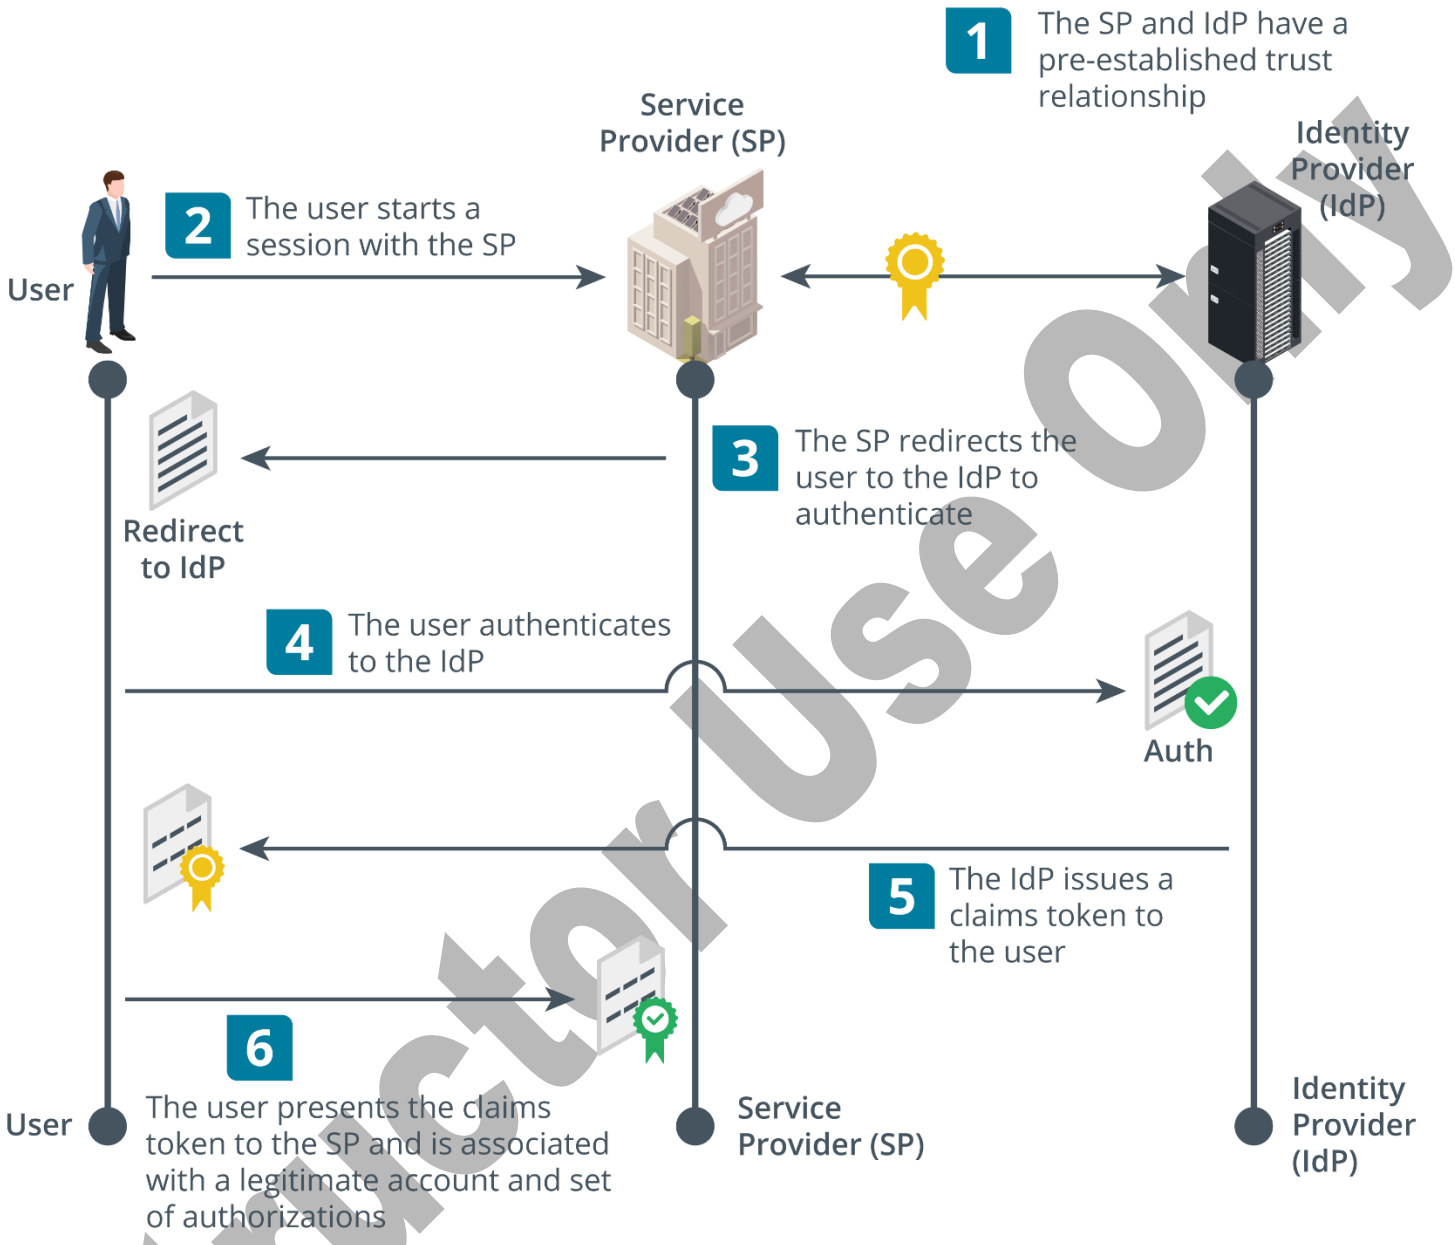

Federated identity management overview. (Images © 123RF.com.)

A federated network or cloud needs specific protocols and technologies to implement user identity assertions and transmit claims between the principal, the relying party, and the identity provider. Security Assertion Markup Language (SAML) is one such solution. SAML assertions (claims) are written in eXtensible Markup Language (XML). Communications are established using HTTP/HTTPS and the Simple Object Access Protocol (SOAP) . The secure tokens are signed using the XML signature specification. The use of a digital signature allows the relying party to trust the identity provider.

An example of a SAML implementation is Amazon Web Services (AWS), which functions as a SAML service provider. This allows companies using AWS to develop cloud applications to manage their customers' user identities and provide them with permissions on AWS without having to create accounts for them on AWS directly.

```
<samlp:Response xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion" ID="200" Version="2.0"
IssueInstant="2020-01-01T20:00:10Z " Destination="https://sp.foo/saml/acs"
InResponseTo="100".
<saml:Issue>https://idp.foo/sso</saml:Issuer>r<saml:Issuer>https://idp.foo/sso</saml:Is
suer>
<ds:Signatur>...</ds:Signature>e<ds:Signature>...</ds:Signature>
<samlp:Statu>...(success)...</samlp:Status.s<samlp:Status>...(success)...</samlp:Status
.
<saml:Assertion xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xs="http://www.w3.org/2001/XMLSchema" ID="2000" Version="2.0"
IssueInstant="2020-01-01T20:00:09Z">
<saml:Issuer>https://idp.foo/sso</saml:Issuer>
<ds:Signature>...</ds:Signature>
<saml:Subjec>...t<saml:Subject>...
<saml:Condition>...s<saml:Conditions>...
<saml:AudienceRestrictio>...n<saml:AudienceRestriction>...
<saml:AuthnStatemen>...t<saml:AuthnStatement>...
<saml:AttributeStatemen>t<saml:AttributeStatement>
<saml:Attribut>...e<saml:Attribute>...
<saml:Attribut>...e<saml:Attribute>...
</saml:AttributeStatemen>t</saml:AttributeStatement>
</saml:Assertio>n</saml:Assertion>
</samlp:Response>
```

Open Authorization

Many public clouds use application programming interfaces (APIs) based on Representational State Transfer (REST) rather than SOAP. These are called RESTful APIs. Where SOAP is a tightly specified protocol, REST is a looser architectural framework. This allows the service provider more choice over implementation elements. Compared to SOAP and SAML, there is better support for mobile apps.

Authentication and authorization for a RESTful API are often implemented using the Open Authorization (OAuth) protocol. OAuth is designed to facilitate the sharing of information (resources) within a user profile between sites. The user creates a password-protected account at an identity provider (IdP). The user can link that identity to an OAuth consumer site without giving the password to the consumer site. A user (resource owner) can grant an OAuth client authorization to access some part of their account. A client in this context is an app or consumer site.

The user account is hosted by one or more resource servers. A resource server is called an API server because it hosts the functions that allow OAuth clients (consumer sites and mobile apps) to access user attributes. An authorization server processes authorization requests. A single authorization server can manage multiple resource servers; equally, the resource and authorization server could be the same server instance.

The client app or service must be registered with the authorization server. As part of this process, the client registers a redirect URL, which is the endpoint that will process authorization tokens. Registration also provides the client with an ID and a secret. The ID can be publicly exposed, but the secret must be kept confidential between the client and the authorization server. When the client application requests authorization, the user approves the authorization server to grant the request using an appropriate method. OAuth supports several grant types—or flows—for use in different contexts, such as server to server or mobile app to server. Depending on the flow type, the client will end up with an access token validated by the authorization server. The client presents the access token to the resource server, which then accepts the request for the resource if the token is valid.

OAuth uses the JavaScript Object Notation (JSON) Web Token (JWT) format for claims data. JWTs can be passed as Base64-encoded strings in URLs and HTTP headers and can be digitally signed for authentication and integrity.

4.2.5 Biometrics and Authentication Technologies (Lesson Video)

Transcript:

In this lesson, we'll look at biometrics and some of the other technologies used for multifactor authentication. Remember, authentication is the process of verifying a claimed identity. And to verify an identity, we need a unique piece of information that could only come from one person.

Biometric authentication has come a long way. You're probably familiar with fingerprint and retina scans, but did you know that people can also be identified by their iris, facial structure, and even by their gait when they walk?
Biometric authentication requires capturing and storing a unique physical attribute in a biometric system. The initial capture is called enrollment. Subsequent authentication attempts are then tested against the stored biometric template.
A few key parameters can determine whether an attribute is viable.
First, the physical attribute should be universal â€“ something that everyone has. Second, the attribute should be unique. It should be measurable and so distinctive that it can be used to tell people apart. Third, the attribute should be permanent â€“ in other words, it should hold up to aging. For example, a person's face might change considerably over time, but their fingerprint or retina will stay the same.
Fourth, the attribute must be collectable. In other words, how easy is it to acquire this measurable attribute? If collecting the measurement is difficult or invasive, it probably won't work. Lastly, the attribute must be difficult to duplicate. For example, if the attribute is the person's voice, it might be possible to evade the system using an audio recording. But if the attribute is the person's retina, then it's much more difficult to circumvent.
If errors do occur, they can be labeled as a false positive or a false negative. A false positive occurs when the biometric system returns a positive match and identifies an unauthorized user as a legitimate user.
A false negative occurs when the biometric system returns a negative match when a legitimate user tries to authenticate.
Now let's look at some other technologies that are used for authentication. You've probably already used a lot of these methods.
To provide an extra layer of security for their employees, some organizations use authenticator applications for their remote employees. An authenticator app, typically installed on a smartphone, provides a 6- to 8-digit code. Every 30 seconds, that code changes. The employee enter their username and password along with the code from the authenticator app, providing additional verification that you are who you say you are.
A similar method that you may have used is a one-time password.
Some banks use this method to allow ATM withdrawals without using a debit card or to let customers access their accounts from the internet.
An application or token creates a one-time password. This password will only work for a single login. After that, the password expires. The password also expires if it isn't used withing a short period of time.
These one-time passwords can be sent to the users several ways.
For example, the one-time password can be sent using the Short Message Service, or SMS. When this option is selected, an SMS message with a one-time code or password is sent to your previously verified phone number. This code can then be used to verify your identify and complete your login.
Similarly, if you chose to receive your code through the phone, you'll get a phone call and an automated voice will provide you with a one-time code or password.
Push notification authentication can also be used to grant access to an account. For example, when you log in to your account, you'll enter your username, but instead of a password, you receive an access request notification on your mobile device. You can then view the authentication details and approve or deny access, typically with a simple press of a button.
One advantage of using this authentication method is that you don't have to remember passwords. However, this method only works with the services provided by a specific company. As a result, there are a limited number of applications that integrate with their services. This makes them unavailable to organizations that use enterprise-level and proprietary solutions.
That's it for this lesson. In this lesson, we looked at biometric systems and several other authentication technologies used for multifactor authentication.

4.2.6 Use a Biometric Scanner (Demo Video)

Transcript:

In this demonstration, we're going to configure biometric authentication systems for Windows. I'm going to use the biometric scanner that's built into the hardware of this laptop, a fingerprint scanner. In order for a computer to use biometrics to authenticate, software to support the login must be on the system. This system has software to handle our fingerprint reader. Each biometric device will need to be supported on the system you wish to use. Based on the type of device you use, you'll configure the operating system to authenticate using the biometrics it offers—your fingerprint, your face, or something else.

Before authenticating with biometrics, you'll need to register with the device. For example, for a fingerprint reader such as the one I'm using, I'll need to register my fingerprint with the device by having it scan the finger a few times. Once registered, that same fingerprint can then be used to authenticate. Authentication happens when the registered biometric, our fingerprint, is compared to the fingerprint that's being scanned on the device. If they match, authentication is granted. Keep in mind that with fingerprints, if your hands are wet or dirty, it will be harder for your biometric system to authenticate you, and your fingerprint will be less likely to match with your previous scan.
Let's demonstrate how to register a fingerprint with this system. I'll search for fingerprint, then click on the Set up fingerprint sign-in option. This opens the Sign-in options Settings page. I'll click on Fingerprint recognition (Windows Hello), then click on Set up.
We're prompted to enter our PIN, a backup authentication method in case the fingerprint isn't available. I've already set up the PIN. But if the PIN hadn't been defined, I would be prompted to configure it. Now, it's asking me to put my finger on the scanner. Then it'll ask me to pick it up and set it down on the reader multiple times while it takes a detailed capture of the fingerprint. This may take a few moments to complete. Now, our configuration is done, and the fingerprint scanner is ready to be used to log on to the system. You also have the option to Add another finger. I only need the one, so that's all I'll add.
And that's all for this demonstration. We talked about how to register your fingerprint with the biometric authentication system.

4.2.7 Use Single Sign-on (Demo Video)

Transcript:

In this demo, we're going to look at single sign-on, or SSO. Now, before we begin, I want to point out that you've probably used single sign-on without even knowing it. In fact, if you have a Gmail or Facebook account, then you've definitely used single sign-on. And if you've ever logged into a Windows domain on a computer, you've used single sign-on there too. Single sign-on is an authentication method that requires you to log in only once to have access to a wide range of services. In a domain, these services include things like file servers and print servers. When you log in, you can access the file server without providing authentication credentials again. It's really nice.

Now, we aren't going to see single sign-on with a domain in this demo. Instead, we're going to look at single sign-on with Google internet services.
Let's start by logging in to a google account. You're probably familiar with Google's email service, Gmail. Let's go to gmail.com and enter our login information. We've logged into Gmail. Because we logged into our single Gmail account, now we have access to all of these Googles services without needing to provide authentication credentials again. There are a lot of them. All of these services fall under the single sign-on of Google.
For example, let's click the YouTube icon, which will take us to YouTube's page. In the top right here, notice that we're automatically signed in. We didn't need to type our username or password. This is because YouTube is connected with Google, which uses a single sign-on service.
The same is true for other Google services. If we go back and click on Google Drive, we have access to that as well. No need to log in.
Let's search LinkedIn in Google. Click Join Now. Notice here, instead of creating an account with LinkedIn directly, we can create an account with Google. This creates a single sign-on link of sorts between LinkedIn and Google that makes it so when we're logged in on Google, we're also logged in on LinkedIn.
Let's try another. Let's go to wordpress.com. Click Log In and look down here. Instead of creating an account, we can link our Google account by clicking Continue with Google. When we do, we're asked which account to link.

We only have one, so I'll click on it. And now, after we wait a minute or two, an account will be created, and we're logged in to the service without needing to provide our password. Again, this is because of the single sign-on functionality that's been implemented between WordPress and Google.
Now, because everything is linked to our Google account, when we log out of it, we're also logged out of those services. Let's do that. We'll log out of Gmail. The page will refresh, and we're asked to log in. So, what happened with Wordpress? Let's go back and look. I'll type â€˜wordpress.com' and press Enter.
Back on WordPress, we're still logged in. This is probably because of how the session cookies are handled. We know we're logged out of Google, so let's log out of WordPress. I'm taken back to the logon page. Click Sign in. We're taken back to the Log in to your account page. I'll click Sign in with Google. Notice my account info here. My password is automatically populated. I'll click Next.
After just a few seconds, I'm logged back into WordPress. Now, guess what else happened? We were actually logging into our Google account again. If I jump back over to my Gmail account, you can see that I was logged back into there, too.
As you can see, single sign-on can be extremely useful and convenient. However, as we know all too well, convenience comes with risk. For example, using single sign-on presents the risk that if someone were to compromise a single account, they could potentially have access to any service that has been linked to that account. One thing that can protect against this is multi-factor authentication, such as requiring a one-time key to access different services.
That's it for this demonstration. In this demo, we looked at single sign-on. We looked at a few different single sign-on features and how a single account can be used to access multiple services.

4.2.8 Biometrics and Authentication Technologies Facts

Biometric authentication is based on a unique physical attribute or characteristic. This type of authentication requires capturing and storing a unique physical attribute with a biometric system.

This lesson covers the following topics:

- Biometric parameters
- Biometric authentication
- Common collection methods

Biometric parameters

For biometric authentication to be a viable security mechanism, it is important to consider the following parameters:

Parameter	Description
Universal	Does each person have the physical attribute being measured?
Unique	Is the physical attribute distinctive enough that it can be used to distinguish between individuals?
Permanent	How well does the specified attribute hold up to aging?
Collectible	How easy is it to acquire this measurable attribute?
Circumvention	Can the attribute be easily circumvented?
Accuracy	Are the results accurate? Accuracy is extremely critical in a biometric system.

Biometric Authentication

The first step in setting up biometric authentication is enrollment:

- A sensor module acquires the biometric sample from the target.
- A feature extraction module creates a template. The template is a mathematical representation of the parts of the sample that uniquely identify the target.

When the user wants to access a resource, they are re-scanned, and the scan is compared to the template. If they match to within a defined degree of tolerance, access is granted.

Biometric authentication can be challenging to implement. The efficacy rate of biometric pattern acquisition and matching and suitability as an authentication mechanism can be evaluated using the following metrics and factors:

- **False rejection rate (FRR)** — is where a legitimate user is not recognized. This is also referred to as a Type I error or false non-match rate (FNMR). FRR is measured as a percentage.
- **False acceptance rate (FAR)** — is where an interloper is accepted (Type II error or false match rate [FMR]). FAR is measured as a percentage.

False rejection causes inconvenience to users, but false acceptance can lead to security breaches, which is usually considered the most important metric.

- **Crossover Error Rate (CER)** — the point at which FRR and FAR meet. The lower the CER, the more efficient and reliable the technology.

Errors are reduced over time by tuning the system. This is typically accomplished by adjusting the sensitivity of the system until CER is reached.

- **Throughput (speed)** — the time required to create a template for each user and the time required to authenticate. This is a major consideration for high-traffic access points, such as airports or railway stations.
- **Failure to Enroll Rate (FER)** — are incidents in which a template cannot be created and matched for a user during enrollment.
- **Cost/Implementation** — some scanner types are more expensive, whereas others are not easy to incorporate on mobile devices.
- Users can find it intrusive and threatening to privacy.
- The technology can be discriminatory or inaccessible to those with disabilities.

Fingerprint recognition is the most widely implemented biometric authentication method. The technology required for scanning and recording fingerprints is relatively inexpensive, and the process is quite straightforward. A fingerprint sensor is usually implemented as a small capacitive cell or optical camera that can detect the unique pattern of ridges making up the pattern. The technology is also nonintrusive and relatively simple to use, although moisture or dirt can prevent readings.

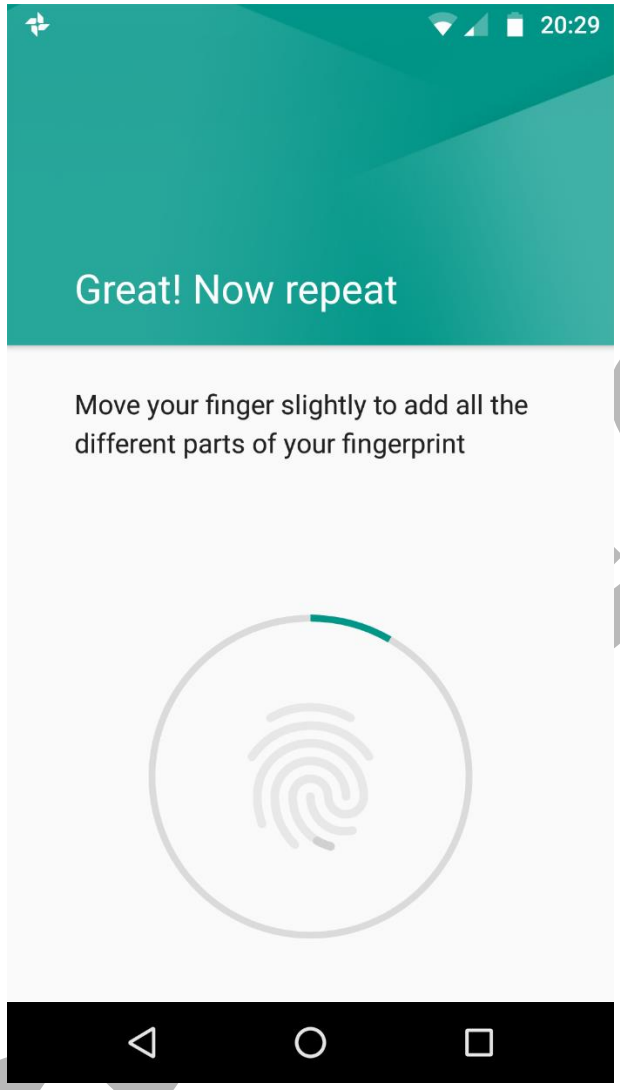

Configuring fingerprint recognition on an Android ™ smartphone. Android is a trademark of Google LLC.

Facial recognition records multiple indicators about the size and shape of the face, like the distance between the eyes or the width and length of the nose. The scan usually uses optical and infrared cameras or sensors to defeat spoofing attempts that substitute a photo for a real face.

Common Collection Methods

Biometric information can be collected for each of the following:

Method	Description
Fingerprints	Fingerprints are made up of patterns of ridges and valleys. Fingerprint scanners analyze these patterns and convert them into a numerical format that can be stored for future comparison.
Retina	The retina is the back portion of the eye sensitive to light. Numerous capillaries move blood to the retina, and these capillaries create a unique pattern. A retinal scanner shines infrared light into an

Method	Description
	eye and measures the amount of reflection. The vessels in the retina absorb infrared light so that the reflection pattern can be stored for future identification.
Iris	The iris is the colorful portion of the eye around the pupil. Infrared light lights up the iris, and the scanner captures images of its unique patterns.
Facial	Facial scanning creates a map of 80 points on an individual's face. The distances measured on this map can be used to identify the person in the future. Measurements could include the distance between the eyes, the nose's shape, the cheekbones' size, etc.
Voice	Voice recognition systems analyze a person's voice for pitch, intensity, and cadence. These systems can be text-dependent or text-independent. Text-dependent authentication requires a specific phrase to be spoken. This could be a pre-determined phrase, or it could be randomly generated. Text-independent authentication uses any speech content.
Vein	Vein recognition scanners use infrared light to determine the vein pattern in your palm. Like a fingerprint, this pattern differs from person to person and does not change. The scanner converts the collected data into a code that is encrypted and assigned to you. The benefits of vein biometrics are: • Veins are internal, so they cannot be altered or covered as easily as hands or a face could be. • More data points can be collected because a palm is larger than an eye or a finger. This provides a higher rate of accuracy. • Internal veins are harder to replicate and can only be captured in proximity.
Gait	Gait recognition analyzes the way that people walk. Each person has a unique way of walking. Several factors determine your gait, including: • Height, weight, and body. • Age. • Health (diseases or disorders). • Personality and emotions. When analyzing gait, the following are measured: • Stride. • Step. • Speed. • Hip and foot angle. • Cadence. Data is gathered using sensors, cameras, or wearable devices. The gait recognition system creates a digital signature that can be stored or compared to existing data. As with most biometric systems, gait recognition systems are still relatively new and should not be used as a stand-alone identification method.

4.2.9 Practice Questions (Section Quiz)

q_authent_biometrics_secp8

The IT security team at a large company is implementing more robust authentication measures to safeguard sensitive data and systems. The team is exploring multi-factor authentication (MFA) options to bolster security. The company deals with highly confidential information and requires a robust solution.

The team has narrowed the choices and is evaluating which aligns BEST with their security needs.

Which multi-factor authentication method utilizes unique physical characteristics of individuals to verify their identity?

Answers:

- ***Biometrics**
- SMS-based one-time passwords (OTP)
- Smart cards
- Passwords and PINs

Explanation:

In biometrics, unique physical characteristics or behavioral traits, such as fingerprints or facial recognition, actively verify a user's identity.

Short message service (SMS)-based one-time passwords (OTPs) actively send temporary codes to users' mobile phones via SMS. While an OTP adds an extra layer of security, it does not actively utilize physical or behavioral traits for verification, making it less secure than biometrics.

Smart cards provide an additional layer of authentication by requiring users to insert or tap a physical card, but they do not involve unique physical or behavioral characteristics like biometrics.

Passwords and personal identification numbers (PINs) actively represent single-factor authentication methods that rely solely on something the user knows.

q_authent_have_secp8

Which of the following are examples of *something you have* authentication controls? (Select two.)

Answers:

- PIN
- Cognitive question
- ***Photo ID**
- ***Smart card**
- Voice recognition
- Handwriting analysis

Explanation:

Something you have authentication controls include physical items that you have on your possession, such as a smart card, photo ID, token device, or swipe card.

Something you know authentication requires you to provide a password, PIN, pass phrase, or the answer to a cognitive question (such as your mother's maiden name).

Something you are authentication uses a biometric system, such as a fingerprint, retina scan, voice recognition, keyboard, or writing recognition.

q_authent_identity_secp8

Which of the following identification and authentication factors are often well known or easily discovered by others on the same network or system?

Answers:

- ***Username**
- Password
- PGP secret key
- Biometric reference profile

Explanation:

The username is typically the least protected identification and authentication factor. Therefore, usernames are often well known or easy to discover, especially by others on the same network or system. The key to maintaining a secure environment is to keep authentication factors secret. Often, usernames are constructed using a standard naming convention, such as first and middle initials plus the full last name, or the first name and last name separated by a period. If these simple construction conventions are known, building usernames from an employee list is very simple.

Passwords, your PGP secret key, and your biometric reference profile are less likely to be well known or easy to discover.

q_authent_know_01_secp8

Which of the following is a password that relates to things that people know, such as a mother's maiden name or a pet's name?

Answers:

- ***Cognitive**
- Dynamic
- One-time
- Passphrase

Explanation:

Cognitive passwords relate to things that people know, such as a mother's maiden name or a pet's name.

Dynamic passwords change upon each consecutive login.

One-time passwords are only valid for a single use.

A passphrase is a password long enough to be a phrase.

q_authent_know_02_secp8

What type of password is *maryhadalittlelamb?*

Answers:

- ***Passphrase**
- Cognitive
- Static
- Composition

Explanation:

A passphrase is a password long enough to be a phrase, such as *maryhadalittlelamb*.

Cognitive passwords relate to things that people know, such as a mother's maiden name or a pet's name,

A static password is created by a user and overseen by an administrator.

Composition passwords are created by the system and are usually two or more unrelated words divided by symbols on the keyboard.

q_authent_know_03_secp8

Which of the following is the MOST common form of authentication?

Answers:

- Photo ID
- Fingerprint
- Digital certificate on a smart card
- ***Password**

Explanation:

Passwords are the most common form of authentication. Most secure systems require only a username and password to provide users with access to the computing environment. Many forms of online intrusion attacks focus on stealing passwords. This makes using strong passwords very important. Without a strong password policy and properly trained users, the reliability of your security system is greatly diminished.

Photo ID, fingerprint, and digital certificate on a smart card are not the most common forms of authentication.

q_authent_multifactor_secp8

Match the authentication factor types on the left with the appropriate authentication factor on the right. Each authentication factor type may be used more than once.

Answers:

- Something you know
- Something you have
- Something you are
- Somewhere you are
- Something you do

Explanation:

Something you know authentication requires you to provide a password or some other data that you know. This is the weakest type of authentication. Examples of something you know authentication controls include:

- Passwords, codes, or IDs
- PINs
- Passphrases (long multi-word passwords)

Something you have (also called token-based authentication) is authentication based on something users have in their possession. Examples of something you have controls include:

- Swipe cards
- Photo IDs
- Smart cards
- Hardware tokens

Something you are authentication uses a biometric system. A biometric system attempts to identify a person based on metrics or a mathematical representation of the subject's biological attribute. Biometric systems are the most expensive and least accepted system type, but are generally considered the most secure form of authentication. Common attributes used for biometric systems include:

- Fingerprints
- Hand topology (side view) or geometry (top-down view)
- Palm scans
- Retina scans
- Iris scans
- Facial scans
- Voice recognition

Somewhere you are authentication (also known as geolocation) is a supplementary authentication factor that uses physical location to verify a user's identity. Examples of implementations include:

- An account is locked unless the user has passed through the building's entrance using an ID card.
- If the user is within RFID range of the workstation, authentication requests are allowed.
- GPS or Wi-Fi triangulation location data is used to determine a device's location. If the user and the device are in a specified location, authentication requests are allowed. If not, the device is locked.

Something you do is a supplementary authentication factor that requires an action to verify a user's identity. Example implementations include:

- Analyzing a user's handwriting sample against a baseline sample before allowing authentication.
- Analyzing a user's typing behaviors against a baseline sample before allowing authentication.

q_authent_passwordless_01_secp8

The IT department at a medium-sized company is exploring ways to enhance its authentication methods to improve security. They want to choose an authentication approach that balances security and user convenience.

Which authentication method eliminates the need for passwords and provides a secure way of verifying a user's identity based on the device's hardware or software characteristics?

Answers:

- ***Passwordless authentication**
- Attestation
- Multi-factor authentication
- Biometric authentication

Explanation:

Passwordless authentication eliminates traditional passwords and relies on other factors such as biometrics, security keys, or mobile push notifications for user verification.

Attestation involves verifying the integrity and authenticity of a device's hardware or software. While it can enhance overall security, it is not a passwordless authentication method.

Multi-factor authentication (MFA) involves using multiple authentication factors, which can include something the user knows (like a password) along with something they have or are (like a fingerprint). While it can enhance security, it still involves using passwords in some cases.

Biometric authentication uses physical characteristics such as fingerprints or facial recognition to verify a user's identity.

q_authent_passwordless_02_secp8

A leading online retail company wants to improve user experience and security for its customers. The security team aims to eliminate the need for users to remember or input complex passwords, reducing the risk of password breaches.

Instead, they propose a solution where users can access their accounts seamlessly through a secure link sent to their verified email or via a push notification on a trusted device. This approach should not involve traditional passwords, fingerprint scans, or multiple validation steps.

Which authentication method is the security team planning to implement for users?

Answers:

- ***Passwordless authentication**
- Attestation
- Multi-factor authentication
- Biometric authentication

Explanation:

Passwordless authentication eliminates traditional passwords and relies on other factors like biometrics, security keys, or mobile push notifications for user verification.

Attestation involves verifying the integrity and authenticity of a device's hardware or software. While it can enhance overall security, it is not a passwordless authentication method.

Multi-factor authentication (MFA) involves using multiple authentication factors, which can include something users know (like a password) along with something they have or are (like a fingerprint). While it can enhance security, it still involves using passwords in some cases.

Biometric authentication uses physical characteristics like fingerprints or facial recognition to verify a user's identity.

q_authent_security_keys_01_secp8

The IT security team at a large tech company is strengthening its authentication methods to protect sensitive company data and systems. The team considered implementing various security measures and understood that each authentication method has distinct features and benefits.

However, they must choose the MOST suitable option that aligns with the organization's security requirements and user convenience.

Which authentication method utilizes a physical device or software to generate secure, unique codes and offers convenience and strong security?

Answers:

- ***Security keys**
- Hard authentication tokens
- Biometric authentication
- Soft authentication tokens

Explanation:

Security keys are authentication devices, either physical hardware or software-based, that generate secure, unique codes for authentication purposes.

Hard authentication tokens are physical devices that generate one-time passwords (OTPs) or passcodes used for authentication. The codes are unique and time-based, so they reduce the risk of unauthorized access.

Biometric authentication relies on unique biological characteristics such as fingerprints, iris patterns, or facial recognition to verify a user's identity. While it offers convenience, it is not explicitly related to security keys.

Soft authentication tokens are software-based and generate one-time passwords (OTPs) or passcodes on a user's device.

q_authent_security_keys_02_secp8

The cybersecurity expert at a technology firm recommends adding another layer of protection to employee accounts. The expert suggests a physical device that users can insert or tap on compatible systems to verify their identity alongside a password.

The proposed solution should use something other than biometric data, produce time-sensitive codes, or be an app or software on personal devices.

Which authentication method is the cybersecurity expert recommending for the employees?

Answers:

- ***Security keys**
- Hard authentication tokens
- Biometric authentication
- Soft authentication tokens

Explanation:

Security keys are authentication devices, either physical hardware or software-based, that generate secure, unique codes for authentication purposes.

Hard authentication tokens are physical devices that generate one-time passwords (OTPs) or passcodes used for authentication. The codes are unique and time-based, so they reduce the risk of unauthorized access.

Biometric authentication relies on unique biological characteristics such as fingerprints, iris patterns, or facial recognition to verify a user's identity. While it offers convenience, it is not explicitly related to security keys.

Soft authentication tokens are software-based and generate one-time passwords (OTPs) or passcodes on a user's device.

q_authent_smart_secp8

A smart card can be used to store all but which of the following items?

Answers:

- Digital signature
- ***Biometric template original**
- Cryptography keys
- Identification codes

Explanation:

A smart card cannot store biometric template originals, as those are physical components of the human body.

A smart card can store digital signatures, cryptography keys, and identification codes.

q_authent_you_are_01_secp8

A technician is assisting a group of new employees with setting up multi-factor authentication.

What philosophy incorporates the use of facial scans or fingerprints to demonstrate authentication?

Answers:

- ***Something you are**
- Somewhere you are
- Something you know
- Something you have

Explanation:

Something you are refers to a biometric or inherence factor. A biometric factor uses physiological identifiers, such as a fingerprint or facial scan, or behavioral identifiers, such as how someone moves (gait).

Somewhere you are means the system applies a location-based factor to an authentication decision. Location-based authentication measures some statistics about where you are.

Something you know means the information used for authentication is from something one can recall, such as a passphrase or username/password combination.

Something you have means the account holder possesses something that no one else does, such as a smart card, key fob, or smartphone that can generate or receive a cryptographic token.

q_authent_you_are_02_secp8

After finding a corporate phone unattended in a local mall, an organization decides to enhance its multi-factor authentication (MFA) procedures.

What MFA philosophy applies a location-based factor for authentication?

Answers:

- ***Somewhere you are**
- Something you know
- Something you have
- Something you are

Explanation:

Somewhere you are means the system applies a location-based factor to an authentication decision. Location-based authentication measures some statistics about where you are.

Something you know means the information used for authentication is from something one can recall, such as a passphrase or username/password combination.

Something you have means the account holder possesses something that no one else does, such as a smart card, key fob, or smartphone that can generate or receive a cryptographic token.

Something you are refers to a biometric or inherence factor. A biometric factor uses physiological identifiers, such as a fingerprint or facial scan, or behavioral identifiers, such as how someone moves (gait).

q_authent_methods_claims-based_sec8

You are a security architect for a large organization that uses various cloud services. The organization wants to implement a system that allows users to authenticate once and then access multiple applications, with the system providing information about the user to the applications.

The system should be able to work across multiple platforms and authentication systems.

Which solution would you recommend?

Answers:

- Local authentication
- ***Claims-based identity**
- Network authentication
- Password-based authentication

Explanation:

Claims-based identity is the correct answer. In a claims-based identity system, an identity provider issues a token containing claims about the user (such as the user's name, role, or privileges) to the user's browser. The user's browser then presents this token to applications, which use the claims to decide what the user is allowed to do. This system can work across multiple platforms and authentication systems.

Local authentication is incorrect because local authentication is used to authenticate a user to a specific system or application. It does not provide a way for users to authenticate once and then access multiple applications.

Network authentication is incorrect because network authentication is used to authenticate a user to a network. It does not provide a way for users to authenticate once and then access multiple applications.

Password-based authentication is incorrect because password-based authentication is a method of authenticating a user based on something the user knows (their password). It does not provide a way for users to authenticate once and then access multiple applications.

q_authent_methods_dn_sec8

As a network administrator, you are tasked with creating a new entry in your company's LDAP directory. You need to ensure that the entry is uniquely identifiable and conforms to the structure of the directory.

Which of the following attributes would you use to accomplish this?

Answers:

- Common name
- Organizational unit (OU)
- ***Distinguished name (DN)**
- Domain component (DC)

Explanation:

Distinguished name (DN) is the correct answer. The distinguished name is a collection of attributes that define a unique identifier for any given resource in an LDAP directory. A DN is made up of attribute-value pairs, separated by commas. The most specific attribute is listed first, and successive attributes become progressively broader.

Common name (CN) is incorrect because the common name is just one attribute of an LDAP entry and is not sufficient to uniquely identify an entry in the directory.

Organizational unit (OU) is incorrect because the organizational unit is just one attribute of an LDAP entry and is not sufficient to uniquely identify an entry in the directory.

Domain component (DC) is incorrect because the domain component is just one attribute of an LDAP entry and is not sufficient to uniquely identify an entry in the directory.

q_authent_methods_federation_01_sec8

Your financial planning company is forming a partnership with a real estate property management company. One of the requirements is that your company open up its directory services to the property management company to create and access user accounts.

Which of the following authentication methods will you be implementing?

Answers:

- ***Federation**
- Single sign-on
- Attestation

- Directory services

Explanation:

In this scenario, you would be implementing a federation authentication method. Federation is the notion that a network needs to be accessible to more than just a well-defined group of employees, such as trusting user accounts created and managed by a different network.

Single sign-on is an authentication process that allows users to access multiple systems, applications, or websites using only a single set of credentials. With SSO, a user authenticates once using designated credentials and can access different resources seamlessly.

Attestation is a protocol used to prove that software can be trusted. It tells the remote user that the application or OS software is legitimate and certified.

Directory services implement single sign-on for resources on the network.

q_authent_methods_federation_02_sec8

You are a network administrator for a multinational corporation that uses various cloud services. The corporation has offices in multiple countries, each with their own local directories for managing accounts and rights.

The CEO wants to implement a system that allows these authorizations to be implemented across all offices and cloud services.

Which solution would you recommend?

Answers:

- Local directory
- Network directory
- ***Federation**
- Single sign-on authentication

Explanation:

Federation is the correct answer. Federation is a system that allows for the implementation of authorizations across different domains, such as multiple offices or cloud services. It uses a federated identity management solution to manage accounts and rights across these different domains.

Local directory is incorrect because a local directory is used to manage accounts and rights within a single office or location. It would not allow for authorizations to be implemented across multiple offices and cloud services.

Network directory is incorrect because a network directory, while it can manage accounts and rights across a network, would not allow for authorizations to be implemented across multiple offices and cloud services, especially if those services are from different vendors.

Single sign-on authentication is incorrect because while single sign-on authentication allows a user to authenticate once and receive authorizations for multiple services, it does not manage accounts and rights across multiple offices and cloud services.

q_authent_methods_kerberos_sec8

You are a network administrator for a large organization that uses a single sign-on (SSO) system for network authentication and authorization.

The organization has recently experienced a security breach, and in response, the CEO wants to implement a system that can authenticate both human users and application services, and that relies on an intermediary to vouch for their identity.

Which protocol would you recommend?

Answers:

- Secure Shell (SSH)
- Lightweight Directory Access Protocol (LDAP)
- Network Time Protocol (NTP)
- ***Kerberos**

Explanation:

Kerberos is the correct answer. Kerberos is a network authentication protocol that works on the basis of tickets to allow nodes communicating over a non-secure network to prove their identity to one another in a secure manner. It can authenticate human users and application services, and relies on a key distribution center to vouch for their identity, which is exactly what the CEO is asking for.

Secure Shell (SSH) is incorrect because SSH is a cryptographic network protocol for operating network services securely over an unsecured network, but it does not provide the intermediary authentication service that the CEO is asking for.

Lightweight Directory Access Protocol (LDAP) is incorrect because LDAP is a protocol used to access and maintain distributed directory information services, but it does not provide the intermediary authentication service that the CEO is asking for.

Network Time Protocol (NTP) is incorrect because NTP is used to synchronize the clocks of computers over a network, but it does not provide the intermediary authentication service that the CEO is asking for.

q_authent_methods_ldap_sec8

You are a network administrator for a large multinational corporation. The corporation has offices in multiple countries and uses various software products from different vendors.

The CEO wants to implement a system that stores information about users, computers, security groups/roles, and services, and allows for interoperability between different vendors' products.

Which directory service would you recommend?

Answers:

- Active Directory
- Novell Directory Services (NDS)
- ***Lightweight Directory Access Protocol (LDAP)**
- X.500

Explanation:

Lightweight Directory Access Protocol (LDAP) is the correct answer. LDAP is a protocol used to access network directory databases, which store information about authorized users and their privileges, as well as other organizational information. Most directory services, including those from different vendors, are based on LDAP, which allows for interoperability.

Active Directory is incorrect because while Active Directory is a directory service, it is a Microsoft product and may not offer the desired level of interoperability with products from different vendors.

Novell Directory Services (NDS) is incorrect because NDS, while a directory service, is a Novell product and may not offer the desired level of interoperability with products from different vendors.

X.500 is incorrect because while X.500 is a standard for directory services, it is not as widely used as LDAP and may not offer the desired level of interoperability with products from different vendors.

q_authent_methods_pam_sec8

You are a system administrator for a company that uses Linux servers. One of your tasks is to implement a new smart-card login system for all employees.

Which Linux authentication method would you use to accomplish this?

Answers:

- Local user account names stored in /etc/passwd
- Password checked against a hash stored in /etc/shadow
- Secure Shell (SSH)
- ***Pluggable Authentication Module (PAM)**

Explanation:

Pluggable Authentication Module (PAM) is the correct answer. PAM is a package for enabling different authentication providers, such as smart-card login. The PAM framework can also be used to implement authentication to network directory services. This would allow you to implement the new smart-card login system for all employees.

Local user account names stored in /etc/passwd is incorrect because this method is used for storing local user account names, not for implementing a smart-card login system.

Password checked against a hash stored in /etc/shadow is incorrect because this method is used for checking a user's password against a stored hash during login, not for implementing a smart-card login system.

Secure Shell (SSH) is incorrect because SSH is used for secure remote login and other secure network services over an insecure network, not for implementing a smart-card login system.

q_authent_methods_remote_sign-in_sec8

A company's IT department has received a request from an employee who is currently working from home. The employee is unable to access the company's internal resources from their home network.

As an IT professional, which type of Windows authentication would you recommend to resolve this issue?

Answers:

- Windows local sign-in

- Windows network sign-in
- ***Remote sign-in**
- Linux authentication

Explanation:

Remote sign-in is the correct answer. Use remote sign-in when the user's device is not directly connected to the local network. Authentication can take place over a virtual private network (VPN), enterprise Wi-Fi, or web portal. These use protocols to create a secure connection between the client machine, the remote access device, and the authentication server. This would allow the employee to access the company's internal resources from their home network.

Windows local sign-in is incorrect because local sign-in is used when the user is physically present at the computer and is logging in directly. It wouldn't help the employee access the company's internal resources from their home network.

Windows network sign-in is incorrect because network sign-in is used when the user is on the same network as the resources they are trying to access. In this case, the employee is not on the company's network, so this wouldn't resolve the issue.

Linux authentication is incorrect because the question specifically asks about Windows authentication. Linux Authentication would not be applicable in this scenario.

q_authent_methods_sso_sec8

Which of the following authentication methods specifically allows users allows users to access multiple systems, applications, or websites using only a single set of credentials?

Answers:

- ***Single sign-on (SSO)**
- Directory services
- Federation
- Attestation

Explanation:

Single sign-on (SSO) allows users to access multiple systems, applications, or websites using only a single set of credentials.

While directory services implements single sign-on for resources on the network, it is not used to refer to single sign-on.

Federation is the notion that a network needs to be accessible to more than just a well-defined group of employees. A company might need to open parts of its network to partners, suppliers, and customers.

Attestation is a protocol used to prove that software can be trusted, telling the remote user that the application or OS software is legitimate and certified.

q_authent_methods_tgt_sec8

You are a network administrator for a company that uses Kerberos for network authentication. An employee is having trouble accessing network resources and you suspect it's an authentication issue.

As part of your troubleshooting process, you decide to check the initial step in the Kerberos single sign-on process.

What should you be looking for?

Answers:

- The principal's request for a service ticket (ST)
- ***The principal's request for a ticket-granting ticket (TGT)**
- The principal's request for a key distribution center (KDC)
- The principal's request for a Secure Shell (SSH)

Explanation:

The principal's request for a ticket-granting ticket (TGT) is the correct answer. The first step in the Kerberos single sign-on process is for the principal (user or service) to request a TGT from the Authentication Service.

The principal's request for a service ticket (ST) is incorrect because the service ticket is not the initial step in the Kerberos process. The service ticket is requested after the TGT has been granted.

The principal's request for a key distribution center (KDC) is incorrect because the principal does not request a KDC. The KDC is a part of the Kerberos system that contains the authentication service and the ticket-granting service.

The principal's request for a Secure Shell (SSH) is incorrect because SSH is a protocol for secure remote login and other secure network services over an insecure network, not a part of the Kerberos authentication process.

q_sso_biometric_parameters_secp8

Which of the following is NOT a parameter considered for biometric authentication?

Answers:

- Universality
- Uniqueness
- Permanence
- Collectability
- ***Color**

Explanation:

The color of the biometric feature does not play a role in the authentication process as it does not contribute to the uniqueness or permanence of the biometric feature.

The following are biometric features:

- Universality refers to the fact that the physical attribute being measured should be present in every individual.
- Uniqueness means that the physical attribute should be distinctive enough to distinguish between individuals.
- Permanence refers to how well the specified attribute holds up to aging.
- Collectability refers to how easy it is to acquire the measurable attribute.

q_sso_crossover_secp8

Which of the following defines the crossover error rate for evaluating biometric systems?

Answers:

- The rate of people who are given access when they should be denied access.
- The rate of people who are denied access when they should be allowed access.
- The number of subjects or authentication attempts that can be validated.
- ***The point where the number of false positives matches the number of false negatives in a biometric system.**

Explanation:

The crossover error rate, or the equal error rate, is the point where the number of false positives matches the number of false negatives in a biometric system.

A false negative (or Type I error) occurs when a person who should be allowed access is denied access.

A false positive (or Type II error) occurs when a person who should be denied access is allowed access.

The processing rate, or system throughput, identifies the number of subjects or authentication attempts that can be validated.

q_sso_facial_recognition_secp8

You are a security manager at a large public event venue. You need to implement a biometric system that can quickly and non-intrusively collect data from a large crowd of people for identification purposes.

Which biometric collection method would be the most suitable for this scenario?

Answers:

- Fingerprint
- Retina
- Iris
- ***Facial**
- Voice

Explanation:

Facial recognition is the most suitable method for this scenario. It can quickly scan multiple faces in a crowd from a distance, making it non-intrusive and efficient for large-scale identification.

While fingerprint recognition is a reliable method, it is not suitable for a large crowd as it requires individual contact with a scanner and can slow down the process significantly.

Retina scanning is highly accurate but requires close contact and can be seen as intrusive. It is not suitable for quickly scanning a large crowd.

Iris recognition is also highly accurate but requires individuals to be close to the scanner and can be seen as intrusive. It is not suitable for quickly scanning a large crowd.

Voice recognition requires quiet environments and can be affected by background noise, making it unsuitable for a large public event venue.

q_sso_failure_to_enroll_secp8

You are the head of security at a large corporation and have recently implemented a new biometric authentication system for access to the company's facilities.

After a few weeks, you notice that a significant number of employees are having trouble registering their biometric data into the system.

Which metric would be most relevant to assess this issue?

Answers:

- False rejection rate (FRR)
- False acceptance rate (FAR)
- Crossover error rate (CER)
- ***Failure to enroll rate (FER)**
- Throughput (speed)

Explanation:

Failure to enroll rate (FER) measures incidents in which a template cannot be created and matched for a user during enrollment. This is the most relevant metric to assess the issue of employees having trouble registering their biometric data into the system.

The following are less relevant to the scenario:

- False rejection rate (FRR) measures the number of valid subjects who are denied access. It would be relevant if legitimate employees were being denied access after successful enrollment, not during the enrollment process.
- False acceptance rate (FAR) measures the number of unauthorized users who are mistakenly allowed access. It would be relevant if unauthorized individuals were gaining access to the facilities, not when employees are having trouble enrolling.
- Crossover error rate (CER) expresses the point at which FAR and FRR meet, with a low value indicating better performance. It would be relevant for assessing the overall performance of the biometric system, not specifically the enrollment issue.
- Throughput (speed) measures the time required to create a template for each user and the time required to authenticate. It would be relevant if the issue was about the speed of the authentication process, not the enrollment process.

q_sso_fingerprint_recognition_secp8

You are a security consultant tasked with implementing a biometric authentication system for a small business. The business owner wants a system that is cost-effective, non-intrusive, and relatively simple for employees to use.

Which biometric authentication method would you recommend?

Answers:

- Retina scanning
- Iris recognition
- Facial recognition
- ***Fingerprint recognition**
- Vein recognition

Explanation:

Fingerprint recognition is cost-effective, non-intrusive, and simple to use, making it the most suitable option for a small business. The technology required for scanning and recording fingerprints is relatively inexpensive and straightforward.

The following are not as ideal in this scenario:

Retina scanning is highly accurate but can be intrusive as it involves shining light into an individual's eye. It's also not as cost-effective as other options, making it less suitable for a small business.

Iris recognition is also highly accurate, it can be costly to implement and may be seen as intrusive by some individuals, making it less ideal for a small business.

Facial recognition can be affected by changes in lighting, the individual's expression, or even their hairstyle. It can also be more costly to implement than fingerprint recognition.

Vein recognition, while highly accurate and secure, requires more expensive technology and can be seen as intrusive, making it less suitable for a small business.

q_sso_gait_recognition_secp8

You are a security consultant for a high-security facility. The facility is looking for a unique, non-intrusive biometric system that can identify individuals even when they are not directly interacting with the system (e.g., not touching a scanner or speaking into a microphone).

Which biometric collection method would be the MOST suitable for this scenario?

Answers:

- Fingerprint
- Retina
- Iris
- Facial
- ***Gait**

Explanation:

Gait recognition analyzes the unique way individuals walk. It can identify individuals from a distance and does not require direct interaction with the system, making it the most suitable option for this scenario.

While fingerprint recognition is reliable, it requires individuals to directly interact with the system by placing their finger on a scanner. It would not meet the facility's requirement for non-direct interaction.

Retina scanning is highly accurate but requires close contact and can be seen as intrusive. It also requires direct interaction with the system, which does not meet the facility's requirements.

Iris recognition is also highly accurate but requires individuals to be close to the scanner. It requires direct interaction and can be seen as intrusive, which does not meet the facility's requirements.

Facial recognition can identify individuals from a distance, but it requires a clear view of the face, which may not always be possible in all scenarios within the facility.

4.3 Authorization

As you study this section, answer the following questions:

- How is authorization different from authentication?
- How does an access control list (ACL) help to increase network security?
- What is the difference between a Discretionary access control list (DACL) and a system access control list (SACL)?

In this section, you will learn to:

- Examine the access token

The key terms for this section include:

Term	Definition
Authorization	Granting a user on the computer system the right to use a resource.
Access control list (ACL)	A collection of access control entries that determines which users are allowed or denied access to an object and the privileges given to that user.
Effective permissions	Access rights are cumulative, giving the user combined permissions from multiple groups.
Deny permissions	Always override Allow permissions.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts.
	• Authentication, Authorization, and Accounting (AAA)
	  ○ Authorization models
	1.4 Explain the importance of using appropriate cryptographic solutions.
	• Obfuscation
	  ○ Tokenization
	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Access control
	  ○ Access control list (ACL)
	  ○ Permissions
	4.6 Given a scenario, implement and maintain identity and access management.
	• Permission assignments and implications

Exam	Objective
	• Single sign-on (SSO) • Access controls ○ Discretionary
TestOut Security Pro	2.5.2 Access control • 2.5.2.1 Access control list (ACL) • 2.5.2.2 Permissions 4.6.8 Access control • 4.6.8.3 Role-based

4.3.1 Authorization (Lesson Video)

Transcript:

In this video, I'll go over authorization. Although you'll probably hear the terms authorization and authentication used interchangeably, they're actually quite different.

First, authentication asks the user who they are and verifies the user's identity so they can access a system. Authorization is what comes after this. It verifies the user's permissions to determine which specific resources the user is allowed access to. Authentication is validated using login credentials, and authorization is validated using an access control list, or ACL.
An ACL is a list of permissions attached to a resource, such as a file or folder. The ACL lists several entries known as access control entries, which are basically users or groups that have access to the resource. Groups allow for efficient administration and make security implementation easier.
A New Technology File System access control list, or NTFS ACL, contains a DACL and an SACL. These are two distinct ACLs types. The D in DACL stands for Discretionary, and the S in SACL stands for System. So Discretionary ACLs specify permissions, and System ACLs are used to keep track of who's accessing a resource and making changes to it.
Another Microsoft term is security principals. Security principals include user accounts, computer accounts, and security group accounts. Each security principal has a unique Security ID called an SID. The system keeps track of all SIDs so that when a user logs onto a Microsoft system, the system creates an access token for him or her.
The access token includes the user account's SID, the SIDs of all security groups the user belongs to, and a list of the user's rights or privileges. User rights are rights that specific users have on the system. These are also known as privileges, and they're defined by the Local Security Policy or through the Microsoft Group Policy. When a user tries to access a resource, such as a shared folder, the user's access token with its SID is compared against the SIDs of the shared folder. If there's a match, the user can gain access to the resource.
For example, let's say that Frank attempts to access the Sales folder. When he does, the system reads the folder's ACL, which contains all users' and groups' SIDs that allow file access. When we look at the ACL here, we see that SID 5 has Read access, SID 7 has Read & write access, and SID 12 has Full control. This means that it's probably an administrator account. When we look at Frank's access token, we see that he has a user SID of 3, and that he's also been assigned to Group1, which has an SID of 7. So when the system compares his token with this folder's ACL, there's one match. Frank has access to the Sales folder because he's a member of Group1. No other SIDs match up, so Frank is only granted Read & write file access. This is a simple example, but you get the idea.
NTFS permissions get a lot more complicated than this. But in essence, they're just a way to allow users access to certain Microsoft system resources.

That's it for this lesson. In this lesson, we discussed authorization. While we mainly focused on Microsoft NTFS systems, most of the terms we used are generic. Access control lists are widely used for assigning user permissions and privileges.

4.3.2 Cumulative Access (Lesson Video)

Transcript:

In this video, I'm going to talk about cumulative access. As a security administrator, one of your most important jobs is to control both file system and network resource access. You do this using permissions, privileges, and roles.

Permissions, privileges, and roles are usually cumulative. As such, it's easy to inadvertently give someone access to something that you don't want them to have access to. When you assign rights to users, groups, and other accounts, you need to ask yourself, "Do I really want this person to have this access level?" If your system supports the concept of a role, you can create a role definition. In some systems, it might called be a role object. This role or role object isn't actually a user account, but we use it in a similar way. A role is kind of like a template. You build a role, add the desired access, and then you can use that shell of an account to create new accounts as needed.
For example, let's say that Sally is the payroll assistant. She's worked at the company for ten years and has access to all the files and systems that she needs to be effective at her job. One day, Sally is promoted to payroll manager, and Frank is hired to take over as payroll assistant. If her user account was created from scratch, you'd have to start all over to redesign Frank's account. But if you'd created Sally's account using a role specifically made for payroll employees, you could use that same role to create Frank's new account. This makes things much easier and much more secure. We simply add or remove users from roles instead of having to remember which user needs which privileges added or removed.
Not every system uses the roles concept. For example, Windows Active Directory doesn't have an object called a role object. But Active Directory groups are more or less the same thing.
Whatever system you're using, you need to remember that each user receives all member groups' access privileges in addition to its own specific assignments. Cumulative access can get out of control pretty quickly. Over time, users end up with too much access to the file system. The best way to reduce this problem is to limit a single user's assigned number of roles. If you can help it at all, don't assign multiple roles.
Take Sally and Frank, for example. Currently, they're both in payroll. But now that Sally is a manager, she needs access to additional employee information, like salaries and financial history. You could solve this by leaving Sally in her current payroll employee role and adding her to the HR role. By doing this, though, you're probably giving her access to too much information. She obviously needs access to salary information, but she doesn't need access to employee reviews or disciplinary files. Instead, you should probably create a payroll management role. If you're dealing with a situation where none of the roles provide the access level needed for a given user, it's a much better option to create a completely new role that has the necessary access level instead of assigning that user to multiple roles.
That's it for this lesson. In this lesson, we talked about cumulative access in relation to roles and groups. We also discussed the importance of limiting a single user's number of roles.

4.3.3 Authorization Facts

Authorization is the process of determining what rights and privileges a particular entity should have on available resources and then enforcing those rights.

This lesson covers the following topics:

- Authorization
- Permissions, privileges, and roles
- Access control lists (ACLs)
- Authorization with single sign-on

Authorization

Authorization means granting the account configured for the user, or the role for a group of users, the rights to use a resource. Authorization manages the privileges granted on resources such as computers, files, and printers. When managing access to resources, be aware of the following:

- A group account is a collection of user accounts that is useful when establishing file permissions and user rights because when many individuals need the same level of access, a group could be established containing all the relevant users.
- When you assign permissions to a group, these permissions are granted to all group members.
- On a Microsoft system, a user right is a privilege or action that can be taken on a system, such as logging on, shutting down, backing up the system, or modifying the system's date and time.
- Permissions apply to objects (files, folders, printers, etc.), while user rights apply to the entire system (the computer).

An authorization model determines how these rights are granted. For example, in a discretionary model, the object owner can allocate rights. In a mandatory model, rights are predetermined by system-enforced rules and cannot be changed by any user within the system.

Permissions, Privileges, and Roles

A crucial part of any security administrator's job is to control access to resources. For example, with file system security, each object in the file system has an access control list (ACL) associated with it. The ACL contains a list of accounts (principals) allowed to access the resource and the permissions they have over it. The order of ACEs in the ACL is essential in determining effective permissions for a given account. ACLs can be enforced by a file system that supports permissions, such as NTFS, ext3/ext4, or ZFS.

Types of permissions are described in the table below.

Permission Type	Description
Effective permissions	Access rights (permissions) are cumulative. If you are a member of two groups with different permissions, you have the combined permissions of both groups (this is known as effective permissions). Effective permissions are the combination of inherited permissions and explicit permissions.
Deny permissions	Deny permissions always override Allow permissions. For example, if a user belongs to two groups and specific permission is allowed for one group and denied for the other, the permission is denied. However, the exception to this rule comes with inherited permissions. If an object has an explicit Allow permission entry, inherited Deny permissions do not prevent access to the object. Explicit permissions override inherited permissions, including Deny permissions.
Cumulative permissions	The following suggestions will help you plan permissions and mitigate issues related to cumulative permissions: • Identify the users and their access needs (the actions each user needs to be able to perform). • Create a group for each type of user with similar needs. Then, make the users members of the appropriate group. • Assign each group (not the user) the permissions appropriate to the group's data access needs. Grant only the necessary permissions. • Take inheritance into account as you assign permissions. Inheritance means permissions granted to a parent container object flow down to child objects

Permission Type	Description
	within the container. Set permissions as high as possible on the parent container and allow each child container to inherit the permissions. • Override inheritance on a case-by-case basis when necessary.

Access Control Lists (ACLs)

Access control lists (ACLs) in computer systems and networks are used to enforce access control policies. An ACL is a list of rules or entries that specify which users or groups are allowed or denied access to specific resources or perform certain actions. In networks, ACLs are associated with routers, firewalls, or similar devices and define rules that determine how network traffic is filtered or forwarded based on criteria like source IP addresses, destination IP addresses, ports, or protocols.

ACLs can help to control network access and protect against unauthorized or malicious activities. ACLs control access to files, directories, or system resources in operating systems and file systems. Each access control entry (ACE) typically contains a user or group identifier and associated permissions controlling actions that are allowed or denied. These permissions often include read, write, execute, and sometimes more granular limits such as modify, delete, or list.

While ACLs offer flexibility and control, managing complex access control policies with numerous ACL entries can become challenging. Complexity increases the risk of misconfigurations. Therefore, proper planning, periodic reviews, and best practice configurations are essential when implementing and maintaining ACLs.

For example, Discretionary access control (DAC) is based on the primacy of the resource owner. In a DAC model, every resource has an owner. The owner creates a file or service, although ownership can be assigned to another user. The owner has full control over the resource, and they can modify its access control list (ACL) to grant rights to others.

An access control system ensures that an information system meets the goals of the CIA triad. Access control governs how subjects/principals may interact with objects. Subjects are people, devices, software processes, or any other system that can request and be granted access to a resource. Objects are the resources. An object could be a network, server, database, app, or file. Subjects are assigned rights or permissions on resources.

A security principal is an object that can be given permissions to an object. Security principals include user accounts, computer accounts, and security group accounts.

- Each security principal is given a unique identification number called a Security ID (SID).
- When a security principal logs on, an access token is generated. The access token controls access to resources and contains the SID for the user or computer, for all groups the user or computer is a member of, and the user rights granted to the security principal.
- When the security principal tries to access a resource or take action, information in the access token is checked. For example, when a user tries to access a file, the access token is checked for the user's SID and all groups. The SIDs are then compared to those in the object's DACL to identify permissions that apply.
- On a Microsoft system, the access token is only generated during authentication. Changes to group memberships or user rights do not occur until the user logs on again and a new access token is created.

Authorization with Single Sign-on

After completing authentication, the client can decrypt the Ticket Granting Service (TGS) session key but not the Ticket Granting Ticket (TGT) . This establishes that the client and key distribution center (KDC) know the same shared secret and that the client cannot interfere with the TGT.

- To access resources within the domain, the principal requests a service ticket (a token that grants access to a target application server). This process of granting service tickets is handled by the TGS.
- The principal sends the TGS a copy of its TGT and the name of the application server it wishes to access, plus an authenticator consisting of a time-stamped client ID encrypted using the TGS session key.

The TGS should be able to decrypt both messages using the KDC's secret key for the first and the TGS session key for the second. This confirms that the request is genuine. It also checks that the ticket has not expired and has not been used before (replay attack). The TGS service responds with the following:

- **A Service session key** — is used between the client and the application server. This is encrypted with the TGS session key.
- **A Service ticket** — contains information about the principal, such as a time stamp, system IP address, Security Identifier (SID) and the SIDs of groups to which it belongs, and the service session key. This is encrypted using the application server's secret key.
- The principal forwards the service ticket, which it cannot decrypt, to the application server and adds another time-stamped authenticator, which is encrypted using the service session key.

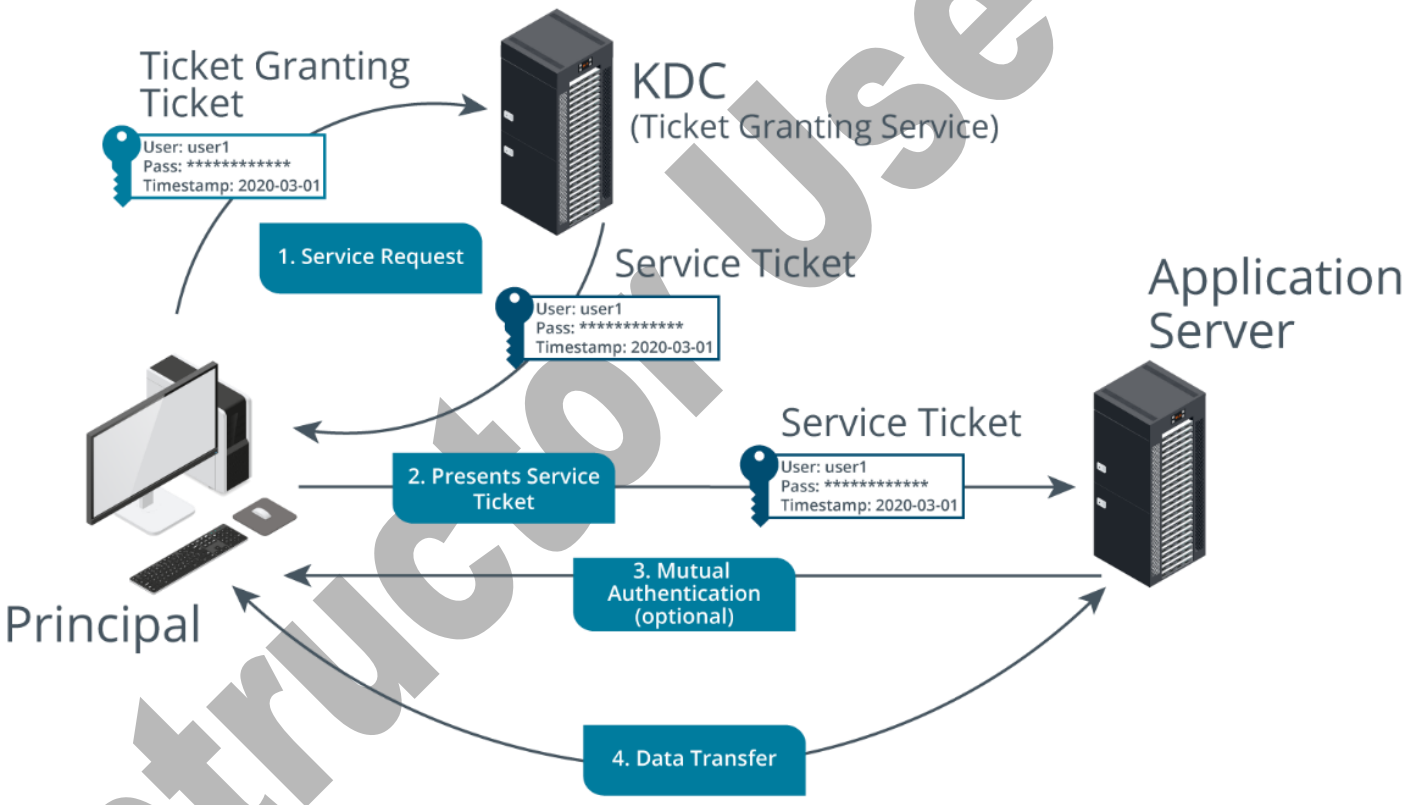

Kerberos Ticket Granting Service. (Images © 123RF.com.)

- The application server decrypts the service ticket to obtain the service session key using its secret key, confirming that the principal has sent it an untampered message. It then decrypts the authenticator using the service session key.
- Optionally, the application server responds to the principal with the time stamp used in the authenticator, which is encrypted using the service session key. The principal decrypts the time stamp, verifies that it matches the value already sent, and concludes that the application server is trustworthy.

This means that the server is authenticated to the principal (referred to as *mutual authentication*). This prevents an on-path attack, where a malicious user could intercept communications between the principal and server.

- The server now responds to access requests (assuming they conform to the server's access control list).

One of the noted drawbacks of Kerberos is that the KDC represents a single point of failure for the network. In practice, backup KDC servers can be implemented (for example, Active Directory supports multiple domain controllers, each of which is running the KDC service).

4.3.4 Examining the Access Token (Demo Video)

Transcript:

In this demonstration, we're going to spend some time working with access tokens. We're going to look at the access token that's granted to a Windows user when they log in to a Windows domain.

I'm currently logged in to this Windows 10 system as the dfellows user in the CorpNet.xyz domain. We can verify this using the â€˜whoami' command at the command prompt. Press Enter, and we see that I'm logged in as dfellows in the CorpNet domain, which is nice to know, but it's not terribly useful. But if we do a â€˜whoami' with the â€˜/all' option, then we see more interesting information. We see the access token that was granted to the dfellows user when I initially logged in.
One of the first things we see is the SID, which is a unique number associated with my user account. In addition, you see my domain and username. We also see the SID of all the different groups that my user account is a member of. First, we see all of the default local groups it's part of. Down here, we also see the domain groups that my account is a member of, such as the Night Shift group, the Day Shift group, and the Enterprise Admins group.
The third section of the access token is Privileges Information. These privileges are also known as user rights. It tells me what my user account can do on the system. For example, there's the Change the system time user right, which my account doesn't have. For me, it's disabled along with a lot of other rights.
Let's go back up here and look at the SID. The key thing to remember about this number, the security identifier, is the fact that when I try to access a file or folder in the NTFS file system, the SID in my access token is compared to the SIDs in the Access Control List of the file or folder. If they match up, I get some form of access, or I may be denied access depending on how the ACL is configured.
At this point, let's go ahead and switch over to my domain controller. On my SRV2019 domain controller, which is the domain controller for my CorpNet.xzy domain, I want to verify the group membership of my dfellows user account that we just looked at.
To do this, I'll come over to Tools and go to Active Directory Users and Computers. In the CorpNet domain, I need to locate my dfellows user account. I happen to know that dfellows is in the CorpNetUsers organizational unit. I'll Select it. Right-click on the account. Let's go to Properties and go to the Member Of tab, and I can see a list of all of the groups that my user is a member of.
I'm a member of Day Shift, Domain Admins, Domain Users, Enterprise Admins, and Night Shift. However, there are many other groups in my domain that I could be a member of. For example, in marketing, I have a group called Marketing. But I'm not currently a member of that group.
To see how the access token determines the level of access I have to files and folders in the NTFS file system, let's come down to File Explorer and go to my C: drive. We have a folder here named SharedFiles, and this folder is currently shared. Let's see who it's shared with. Click on Properties. Go to Sharing. Let's go to Advanced Sharing and then to Permissions.
Here, you can see that two groups have been granted access to this folder. These are the groups that are in the ACL of the SharedFiles folder. Anyone who's a member of either of these groups receives the permissions defined below.
We want to see how the access token works, so let's remove these groups--remember, my dfellows user account is a member of both of these. Let's go ahead and remove these groups and add another group that my user account isn't a member of, for example, the Marketing group we just looked at.
Let's grant Marketing full control. Let's remove Everyone, click Remove, and then click Administrators. Click Remove once again. Hit Apply, OK, OK again, and Close. The SharedFiles folder is now shared with any user who's a member of Marketing. dfellows is not a member of Marketing. Let's go ahead and see what happens now if I try to access this folder from my client system.

Back on my client system, where I'm still logged in as the dfellows user, I come down to Search, and let's enter â€˜\\srv2019'. Press Enter, and it should open File Explorer with a list of all the shares currently available on the file server. There's our SharedFiles share.
Within this share, there are a couple of documents. Let's see if I can access them. If I double-click, there's a problem. I can't access the folder because of the Access Control list. My access token isn't there. I'm not on the list, so I can't access the share. This is an example of Implicit Deny. Remember, with Implicit Deny, if there's no match, then, automatically, access is denied. That's the case here. Only a user who's a member of the Marketing group can access this share.
Let's go ahead and fix this. I'll close this window.
Let's go back over to the domain controller and add the dfellows user to the Marketing group. Here's my Marketing group that I want to add my user account to. I can do this in two ways. I can right-click here, go to Properties, and then go to the Members tab and add my user account there. Or I could come to the user account itself, right-click, and say Add to Group. Either way, it does the same thing.
Let's add it to the Marketing group. Group operation was successful. Click OK. Let's just verify that. Go to Properties. Go to the Member Of tab. Sure enough, I'm a member of Marketing now. OK.
Now let's go back over to my Windows 10 workstation and try accessing the share again. Let's come back down to Search, and let's access our list of shares available. There's the SharedFiles share. If I double-click on it, I still can't access it. Why not? I was added to the Marketing group. The Marketing group is in the Access Control list of the share. In fact, it has full control over the share. Why can't I access it?
The key thing to remember here is that the access token over here, for my dfellows user account, was created when I logged in. I made the change to my group membership for my user account after I had already logged in. In fact, if we come down to the bottom and we type the â€˜whoami /all' command again, you can see that I'm still a member of just Research, Support, and Enterprise Admins as far as my access token is concerned. It doesn't list the Marketing group. That's because we've made the change in group membership since I logged in.
In order to fix this and add the Marketing group to the access token, I have to log off as dfellows on the system and then log back in as that user account again, which will re-issue my access token. This should make this account a member of the Marketing group and give us access to the shared folder.
Let's go ahead and do that. All right. I've logged off, and now I've logged back in. Let's check my access token to see if the Marketing group has been added. Do â€˜whoami /all' one more time. Notice, when I scroll down here and look, I can see that the Marketing group has been added.
Let's test it. We'll come down here, open up my run box again, and browse the various shares available on the server. There's my ShardFiles share. If I double-click on it this time, I have access to the files. Remember, we gave the Marketing group full control, so I'm actually allowed to go in and have access to the file. I can add, modify, delete, or do anything that someone with full access can do. I'm able to do that because now, the Marketing group is in the access token of the dfellows user account.
That's it for this demonstration. In this demo, we looked at the access token that's assigned to a Windows user when they log in to the Windows domain. We've looked at the group membership contained within the access token. Then we did a little experiment to show how the access token is assigned when the user logs in, and any changes made to the user account's group membership aren't reflected until the next login.

4.3.5 Implement an Access Control Model

4.3.6 Practice Questions (Section Quiz)

q_authorize_acl_secp8

Which security mechanism uses a unique list that meets the following specifications:

- The list is embedded directly in the object itself.
- The list defines which subjects have access to certain objects.
- The list specifies the level or type of access allowed to certain objects.

Answers:

- Mandatory access control
- Conditional access
- Hashing
- ***User ACL**

Explanation:

A user ACL (user access control list) is a security mechanism that defines which subjects have access to certain objects and the level or type of access allowed. This security mechanism is unique for each object and embedded directly in the object itself.

Mandatory access control (MAC) is an access control system based on classifications of subjects and objects to define and control access.

Conditional access is a way to enforce access control while also encouraging users to be productive wherever they are.

Hashing is a cryptographic tool that creates an identification code that is employed to detect changes in data.

q_authorize_authorize_secp8

What is the process of controlling access to resources such as computers, files, or printers called?

Answers:

- Mandatory access control
- Conditional access
- ***Authorization**
- Authentication

Explanation:

Authorization is the process of controlling access to resources such as computers, files, or printers.

Mandatory access control (MAC) is an access control system based on classifications of subjects and objects to define and control access.

Conditional access is a way to enforce access control while also encouraging users to be productive wherever they are.

Authentication is the verification of the issued identification credentials.

q_authorize_group_secp8

Which of the following objects identifies a set of users with similar access needs?

Answers:

- ***Group**
- DACL
- SACL
- Permissions

Explanation:

A group is an object that identifies a set of users with similar access needs. Microsoft systems have two kinds of groups, distribution groups and security groups. Only security groups can be used for controlling access to objects.

A discretionary access control list (DACL) is an implementation of discretionary access control (DAC).

A system access control list (SACL) is used by Microsoft for auditing in order to identify past actions performed by users on an object.

Permissions define the rights and access users and groups have with objects.

q_authorize_permission_secp8

Which of the following identifies the type of access that is allowed or denied for an object?

Answers:

- DACL
- User rights
- SACL
- ***Permissions**

Explanation:

Permissions define the rights and access users and groups have with objects. Permissions are applied to objects such as files and folders.

A discretionary access control list (DACL) is an implementation of discretionary access control (DAC).

On a Microsoft system, a user right is a privilege or action that can be taken on a system, such as logging on, shutting down, backing up the system, or modifying the system date and time.

A system access control list (SACL) is used by Microsoft for auditing in order to identify past actions performed by users on an object.

q_authorize_sacl_secp8

Which of the following is used by Microsoft for auditing in order to identify past actions performed by users on an object?

Answers:

- DACL
- User rights
- ***SACL**
- Permissions

Explanation:

A system access control list (SACL) is used by Microsoft for auditing in order to identify past actions performed by users on an object.

A discretionary access control list (DACL) is an implementation of discretionary access control (DAC).

On a Microsoft system, a user right is a privilege or action that can be taken on a system, such as logging on, shutting down, backing up the system, or modifying the system date and time.

Permissions define the rights and access users and groups have with objects. Permissions are applied to objects such as files and folders.

q_authorize_security_secp8

Which type of group can be used for controlling access to objects?

Answers:

- ***Security**
- DACL
- Distribution
- Authorization

Explanation:

Only security groups can be used for controlling access to objects.

A discretionary access control list (DACL) is an implementation of discretionary access control (DAC).

Distribution groups cannot be used for controlling access to objects.

Authorization is the process of controlling access to resources such as computers, files, or printers.

q_authorize_service_ticket_secp8

You are a security administrator for a large organization that uses Kerberos for authentication.

One day, you receive a call from a user who is having trouble accessing a specific service on the network. You suspect that the issue might be related to the Kerberos service ticket.

Which of the following steps should you take first to troubleshoot this issue?

Answers:

- Invalidate all active service tickets for the user.
- ***Check the expiration time of the user's service ticket.**
- Request a new Ticket Granting Ticket (TGT) for the user.
- Change the user's password in the Kerberos database.

Explanation:

The first step in troubleshooting should be to check the expiration time of the user's service ticket. Kerberos service tickets have a limited lifespan for security reasons. If the service ticket has expired, the user will need to request a new one.

Invalidating all active service tickets for the user is a drastic step that could disrupt the user's access to other services. This should not be the first step in troubleshooting.

Requesting a new Ticket Granting Ticket (TGT) for the user is not necessary if the problem is specifically with the service ticket. The TGT is used to obtain service tickets, but it is not directly involved in accessing services.

Changing the user's password in the Kerberos database is a drastic step that could disrupt the user's access to the network. This should not be the first step in troubleshooting, and it is not directly related to the service ticket.

q_authorize_tgt_secp8

You are the IT security manager at a large organization that is implementing a single sign-on (SSO) solution for the first time. The SSO solution uses the Kerberos protocol.

During a meeting, your team discusses the following options for the initial step in the Kerberos authentication process.

Which option should be the initial step in the Kerberos authentication process?

Answers:

- The client sends a request to the Ticket Granting Service (TGS) for a service ticket.
- ***The client sends a request to the Authentication Server (AS) for a Ticket Granting Ticket (TGT).**
- The client sends a request to the service server for a service ticket.
- The client sends a request to the Key Distribution Center (KDC) for a session key.

Explanation:

The first step in the Kerberos authentication process is for the client to send a request to the Authentication Server (AS) for a Ticket Granting Ticket (TGT). The AS verifies the client's credentials and, if valid, issues a TGT. The TGT is then used to request service tickets from the TGS. This is the correct answer.

The client does not initially send a request to the Ticket Granting Service (TGS) for a service ticket. The TGS is involved later in the process, after the client has received a TGT from the AS.

The client does not initially send a request to the service server for a service ticket. The service server is involved later in the process, after the client has received a service ticket from the TGS.

While the Key Distribution Center (KDC) is a critical component of the Kerberos protocol, the client does not directly request a session key from the KDC. The KDC houses both the AS and the TGS, and it is the AS that initially provides the TGT.

q_authorize_token_01_secp8

Marcus White has just been promoted to a manager. To give him access to the files that he needs, you make his user account a member of the Managers group, which has access to a special shared folder.

Later that afternoon, Marcus tells you that he is still unable to access the files reserved for the Managers group.

What should you do?

Answers:

- ***Have Marcus log off and log back in.**
- Manually refresh Group Policy settings on his computer.
- Add his user account to the ACL for the shared folder.
- Manually refresh Group Policy settings on the file server.

Explanation:

On a Microsoft system, an access token is only generated during authentication. Changes made to group memberships or user rights do not take effect until the user logs in again and a new access token is created.

Use NTFS and share permissions, not Group Policy, to control access to files. In addition, Group Policy is periodically refreshed, and new settings are applied on a regular basis.

q_authorize_token_02_secp8

Which of the following terms describes the component that is generated following authentication and is used to gain access to resources following login?

Answers:

- ***Access token**
- Account policy
- Cookie
- Proxy

Explanation:

When a security principal logs on, an access token is generated. The access token is used to control access to resources and contains the following information:

- The security identifier (SID) for the user or computer
- The SID for all groups the user or computer is a member of
- User rights granted to the security principal

When the security principal tries to access a resource or take an action, information in the access token is checked. For example, when a user tries to access a file, the access token is checked for the SID of the user and all groups. The SIDs are then compared to the SIDs in the object's DACL to identify permissions that apply.

Account policies in Group Policy control requirements for passwords, such as minimum length and expiration times.

Cookies are text files that are stored on a computer to save information about your preferences, browser settings, and web page preferences. Cookies identify you (or your browser) to websites.

A proxy is a server that stands between a client and destination servers.

q_authorize_token_03_secp8

Lori, who has been a member of the Project Management group, was recently promoted to manager of the team. She has been added as a member of the Managers group.

Several days after being promoted, Lori needs to have performance reviews with the team she manages. However, she cannot access the performance management system. As a member of the Managers group, she should have the Allow permission to access this system.

What is MOST likely preventing her from accessing this system?

Answers:

- ***She is still a member of the Project Management group, which has been denied permission to this system. Deny permissions always override Allow permissions.**
- She is still a member of the Project Management group, which has been denied permission to this system. However, being a member of the Managers group should allow her to access this system. Allow permissions always override Deny permissions. There must be an explicit permission entry that is preventing her from accessing the management system.
- Her user object has been assigned an explicit Deny permission to the performance management system.
- Her user object has been assigned an explicit Allow permission to the performance management system, but she inherited the Deny permission assigned to the Project Management group (which she still belongs to). Inherited Deny permissions override explicit Allow permissions.

Explanation:

The most likely cause of this problem is that Lori is still a member of the Project Management group, which has been denied permission to this system. Deny permissions always override Allow permissions.

Allow permissions do not override Deny permissions unless the Allow permission is explicitly assigned and the Deny permission is inherited. It is unlikely that her user object has been assigned an explicit Deny permission to the performance management system since best practice is to assign permissions to groups, not to users.

q_authorize_user_secp8

Which of the following is a privilege or action that can be taken on a system?

Answers:

- DACL
- ***User rights**
- SACL
- Permissions

Explanation:

On a Microsoft system, a user right is a privilege or action that can be taken on a system: such as logging on, shutting down, backing up, or modifying the date and time. User rights apply to the entire system.

A discretionary access control list (DACL) is an implementation of discretionary access control (DAC).

Microsoft uses a system access control list (SACL) for auditing in order to identify past actions performed by users on an object.

Permissions define the rights and access users and groups have with objects. Permissions are applied to objects such as files and folders.

4.4 Active Directory Overview

As you study this section, answer the following questions:

- What is the purpose of a domain?
- How do organizational units (OUs) simplify security administration?
- How do computer policies differ from user policies?
- What is the order in which Group Policy objects (GPOs) are applied?

In this section, you will learn to:

- Join a domain
- Manage Active Directory objects
- Create OUs
- Delete OUs
- Use Group Policy
- Create and link a GPO
- Create user accounts
- Manage user accounts
- Create a group
- Create Global Groups

The key terms for this section include:

Term	Definition
Domain	A domain is an administratively defined collection of network resources that share a common directory database and security policies. The domain is the basic administrative unit of an Active Directory structure.
Tree	A tree is a group of related domains that share the same contiguous DNS namespace.
Forest	A forest is a collection of related domain trees. The forest establishes the relationship between trees that have different DNS namespaces.
Organizational unit (OU)	An organizational unit is similar to a folder. It subdivides and organizes network resources within a domain.
Object	Each resource within Active Directory is identified as an object.
Domain controller	A domain controller is a server that holds a copy of the Active Directory database. It is also the copy of the Active Directory database on a domain controller that can be written to.
Replication	Replication is the process of copying changes to Active Directory on the domain controllers.
Member servers	Member servers are servers in the domain that do not have the Active Directory database.
Policy	A policy is a set of configuration settings applied to users or computers.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts. • Authentication, Authorization, and Accounting (AAA)

Exam	Objective
	4.5 Given a scenario, modify enterprise capabilities to enhance security. • Operating system security ○ Group Policy
TestOut Security Pro	1.1 Manage identity 1.1.1 Manage Windows local and domain users and groups 1.1.3 Manage Active Directory OUs

4.4.1 Active Directory Introduction (Lesson Video)

Transcript:

In this lesson, we're going to introduce you to the basics of Microsoft's Active Directory service, or AD service. Using Active Directory can be beneficial when managing larger computer environments.

In standalone and workgroup networking models, all computer management takes place on each individual computer. This includes authentication and security.
This means that every computer stores its own authentication independently from any other computer.
Although these models may work well in small environments of 10 to 15 computers, they don't work well in larger environments. This is because duplicate sets of user accounts must be created and maintained on each computer.
To overcome these types of challenges, Microsoft created Active Directory services.
AD is a centralized database that provides three major benefits: centralized resources and security administration, single logon for access to global resources, and simplified resource location.
Let's first discuss centralized resources and security administration.
This means that an administrator can manage and secure their network resources and security objects from a single point. That single point is the Active Directory database.
This means thatâ€”unlike the workgroup networking modelâ€”they don't have to go from computer to computer to implement these features.
This feature is probably the most compelling reason for using Active Directory.
The second benefit is that you can provide access to the network resources on any server in the domain with a single logon.
So once a user has been created on a domain controller, that user can be granted access to any resource in the networkâ€”from any network computer.
For example, if this computer is this user's main desktop computer, he or she could authenticate or log in to this computer daily. He or she would never have to worry about restricted access.
However, if the user needed this computer to attend a meeting in a conference room, they could also login or authenticate from it. They would still have access to the same resources they had when logged in from the main computer.
The third benefit is the ability to simplify access to files and print resources.
This is done by allowing files and print resources to be published on the network. Once done, a user can search the Active Directory database for the desired resource and then securely access it.
For example, if our user authenticated the domain from this computer but wanted to access a shared folder located on this computer, they wouldn't need to directly reauthenticate to that computer.
Instead, they could simply access the folder for which they had been granted rights.
Likewise, if our user needed to print files, they could find and print to the network printer without needing to reauthenticate to the printer.
As you can imagine, it's going to take a little more work to implement these features than would be required to manage a workgroup network.

As such, you need to be familiar with the various components that comprise an Active Directory implementation.
Let's look at what they are.
The first thing we need to look at is the concept of a domain.
A domain is an administratively defined collection of network objects or resources that share a common directory database and security policies.
The domain is the basic administrative unit of Active Directory. It's the Active Directory database.
As such, it stores our user information and our security information.
For example, in this network, you can see we're working with the CorpNet.xyz domain.
Notice that the domain is identified using a distinguished Domain Name System name, or DNS name, of CorpNet.xyz.
But sometimes we refer to a domain by just its domain or common name. In this case it's CorpNet.
Another component of Active Directory is the domain controller.
A domain controller, also called a DC, is a Windows server that holds a copy of the Active Directory database.
A single domain can have multiple domain controllers, with each domain controller containing a copy of the AD database.
In addition, any changes that need to be made to the Active Directory database can be made to any of the domain controllers. Then that change will be sent out to all the other domain controllers within the domain.
This process is called replication.
For example, if I create a new user account on this domain controllerâ€" through the process of replicationâ€"that same user account would be copied or replicated to the Active Directory databases on all the other domain controllers in the domain.
This allows Active Directory to maintain consistency between all the domain controllers. No matter which one you're working with, it's going to have the exact same information as all the other ones in that domain.
Next, we need to look at organizational units, or OUs.
An OU is like a folder that's used to subdivide and organize resources within the domain.
Examples of organizational units are shown here. We have Marketing, Sales, and Research.
An organizational unit is also known as a container object.
This means that it can contain other domain objects and can even contain other organizational units. Take advantage of this if you want to nest multiple OUs within a given organization unit.
For example, when I expand the Sales OU, you see that I've used several OUs to organize my sales team by region.
OUs can also contain other types of objects, such as user objects, computer objects, and so on.
Besides logically organizing your network resources, OUs also help you simplify your security administration.
For example, if all my users in the Research Department need a similar level of access to network resources, then I can apply the necessary security policies only to the Research OU. Every user will receive those same policies.
This eliminates the need to go through each Research OU user and make the same security changes over and over.
It also means that if I add a new user to this OU, that user will automatically receive the same security privileges. They don't need me to manually assign them anything.
Likewise, if I were to delete a user from the Research OU, that user would lose whatever security privileges they received in the first place.
Active Directory also includes a special type of container object called a built-in container.
A built-in container looks like an organizational unit, but it's not.
Some of the built-in containers shown here include Foreign Security Principles, Managed Service Accounts, and Users.
Notice that the icons we use for built-in containers are slightly different than those used for the OUs we looked at just a minute ago.
Like the other type of organizational units, built-in containers can contain other objects and can be used to organize our Active Directory objects.
But, these built-in containers are created by default when the domain is initially installed. They're different from organizational units in several key ways.
For example, you can't create, move, rename, or even delete a built-in container.
In addition, they have very few properties that can be edited. There's also not much I can configure with one of these.
With organizational units, on the other hand, there's much more that can be configured.
In addition to organizational units and built-in containers, Active Directory also includes objects.
With Active Directory, each network resource is identified as an object.
For example, in the Research OU, you see I have four objects. Three user objectsâ€"Pascal, Paul, and Sailorâ€"and a group object named USA.
And there are many other object types used within Active Directory as well.
You can have objects for computers, for printers, and more.
Each of these objects contains attributes. Each attribute represents some bit of information about that object.

Let's suppose that we're viewing the attributes of the Sailor Olsen user object.
Some of the attributes associated with a user object are the user's name, phone number, email address, and so on.
These attributes are used for both locating and securing network resources.
Essentially, if you want someone to be able to log in to the domain and access the network resources, then you must create a user object for them within the domain.
That's it for this lesson.
In this video, we introduced you to Active Directory.
We first talked about Active Directory domains.
We talked about domain controllers.
We talked about organizational units.
We talked about built-in containers.
Then we talked about individual objects in Active Directory.

4.4.2 Joining a Domain (Demo Video)

Transcript:

In this demonstration, we're going to review how to join a Windows workstation to an Active Directory domain.

I'll come down to the Start button, right-click, and then go to System. Beside Related links, click on Domain or workgroup. This will launch our System Properties window. Toward the bottom of the window, you can see that it says, "To rename the computer or change to a domain or workgroup, click Change." That seems straightforward, so let's click on Change.
Down here, you can see it's part of a workgroup, and the name of the workgroup is WorkGroup. This is the default workgroup that's set up when you install a Windows system. It's not joined to a domain yet. We need to change that.
In order to join this system to a domain, we need to create a domain. That means we have to install a Windows server system, install Active Directory on it, make it a domain controller, and then configure a domain. Let's switch over to my Windows Server system.
As you can see here, on this server system, I've already done all of those tasks. This is my Windows server system. Note, over here, that it already has the Active Directory Domain Services role installed. DNS is installed, and the server is functioning as the domain controller for the domain named CorpNet.xyz.
Before we switch back over to the workstation system, I want to point out something very important: this server is also functioning as my DNS server for the domain. Depending on the implementation you're working in, this may or may not be the case in the real world. I've set it up this way for convenience in a demonstration environment. I have all of my services, including directory services and DNS, all running on the same system.
In a large implementation, it's very likely that the DNS server would be configured to run on a different server. We don't have to worry about that here. We just need to remember that this server, with an IP address of 10.10.10.5, is my DNS server for my CorpNet domain.
With that information in mind, let's go ahead and join this workstation to the CorpNet.xyz domain.
We need to get back to the System Properties window, where we were a few minutes ago. Keep in mind that there are a few different ways to get to System Properties, and as Windows evolves, it could be different in the future.
On the Computer Name tab, we have an option called Change. Click it. Notice we could change the computer name, here, if we wanted to. This is also where we can make this workstation a member of a domain.
I'm not going to do that just yet because currently, this workstation doesn't know how to contact the domain controller. It's not configured to use my Active Directory DNS server. We need to fix that before we can proceed.
Let's go down to our Network icon, down here in the notification area. Right-click and go to Network and Internet settings. We'll click here, where it says Ethernet. If I scroll down a bit, you can see our IP assignment, which is set to Automatic, or DHCP. Our IP is coming via DHCP, and so is our DNS information. And in this case, that's a problem Let's fix it.
If your network is set up so that your DHCP server is handing out the correct IP address for your Active Directory DNS server, you don't actually have to make any changes. If it's not, then you need to go manually specify the IP address pointing to the DNS server that's functioning within your Active Directory environment. That's the case we have here. I happen to know that the DHCP server on this network is handing out a different IP address for the DNS server than the one that's being used by my Active Directory domain.

Before I'll be able to contact the domain controller and locate the CorpNet.xyz domain, we're going to need to change that. Let's click on Edit, click the dropdown, and change this setting to Manual. Now we'll change this one from Off to On. I could just change the DNS information on this workstation, but I'm going to go ahead and configure it with a static IP address. That's just my preference. In other words, I could manually set my DNS to point to my Active Directory server and obtain the IP from the DHCP dynamically.
Let's configure the settings now. For the IP, I'll use 10.10.10.199. I know this isn't an IP that my DHCP server is handing out, so I'm safe to use it here. My subnet mask is 255.255.255.0. The gateway is my router IP address, which is 10.10.10.1. Now the DNS info, 10.10.10.5.
Remember, I pointed out a minute ago that it's very important to remember that DNS is running on the same server as Active Directory. Its IP address is 10.10.10.5. So that's the value we just put into this field. We'll click Save.
I'm going to come down to Search, type in cmd to get to a Command Prompt, use elevated privileges, and type Ipconfig /all. Note that my we can confirm that our DNS server has been changed to 10.10.10.5. That's good. All is well in the world. Now my workstation will be able to contact my domain controller in the CorpNet.xyz domain.
With that set up, we're ready to change settings and join the domain. We'll go back to System Properties. I'll mark Domain and then enter in the name of the domain that I want to join, CorpNet.xyz, and click OK.
Of course, before I can join this workstation to the domain, I have to provide the credentials of an administrative domain user. Not the local system, but an administrative user in the domain itself. That username is dfellows.
As I said earlier, we're not looking at an administrator user on this local system; we're looking at for an administrator in CorpNet.xyz. So I have to be sure to put in the right password for that user account. I'll go ahead and enter in the password for the administrator user account in my CorpNet.xyz domain. Click OK.
Everything went well. I was able to authenticate. I know that because it's welcoming me to the domain. I'll click OK. It does require the system to restart before the changes can take place. I'll click OK, and I'll Restart Now. We'll click that option and wait just a minute while the system reboots.
At this point, my system has rebooted. I have the option of logging in. I was logged in as my local user prior to joining the domain, and I can still log in as my local user after joining the domain as well. But now there's a new option available: Log in as Other user. I can now log in using a domain user account instead of a local user account.
You see, domain user accounts have a lot of advantages. For example, I could log in to this workstation using a domain user account. I could go to other workstations in the domain and log in using that very same user account. This means I don't have to duplicate my user account on every single workstation in my network that I'm going to want to log in from. If I were functioning in a workgroup situation, I would. I would have to go through and make myself a local user on each and every workstation in my network and configure that user account with the same password. That would be a pain. But by using a domain user account, I can have one user account defined in the domain and use it to log in to all of the workstations in the domain.
If you're going to log in as a domain user account, you need to log in a little bit differently. You'll note down here that it tells me that I'm going to be signing into CorpNet. So, technically, I could type dfellows right here, and it would authenticate me using the dfellows user account in the domain instead of the local user account on the local system. However, there may be times when this option is not automatically configured, or you need to sign in to another domain that's not listed here. The way you do that is to type the name of the domain first. In this example, we're going to log in to CorpNet. Then I put a backslash, and then I put in the name of the user in the domain that I want to use for authentication.
In this case, I'm going to log in as the dfellows user in the CorpNet domain. I'm essentially telling the workstation where to locate the user account I want to log in from. Because I specified that the dfellows user here resides in CorpNet, it won't use the local user account on the local system to authenticate. Instead, it'll use the dfellows user in the domain. I'll go ahead and log in with my password.
Because I'm logging in to the system for the first time using a domain user account, we need to provision that user account on this system. I have to have my profile directory set up. And I need to have my default apps installed. All of this has to happen for that domain dfellows user account on the local system. It'll take just a minute to complete.
The provisioning process is complete, so my dfellows domain user now has a profile on this local workstation.
That's it for this demonstration. In this demo, we joined a Windows workstation to a domain.

4.4.3 Managing Active Directory Objects (Demo Video)

Transcript:

In this demonstration, we're going to spend some time managing Active Directory objects.

Notice that we're going to perform these tasks from a Windows server system. I'm using the Server Manager utility on the desktop of a Windows server system. It's named DC1. This particular server system has the Active Directory Domain Services role installed, making it a domain controller. It's hosting the CorpNet.xyz domain.

To manage the objects within this domain, within Server Manager, I click Tools and then click Active Directory Users and Computers. When I do, the Active Directory Users and Computers utility is displayed.

Before we go any further, be aware that you don't have to run this utility right off of the server desktop. If you want to, you can install the necessary software, such as Windows Admin Center, on a workstation. That's beyond the scope of this demo, so we're just going to perform these tasks from the server, which is totally fine for what we're doing today.

Notice that when I initially load Active Directory Users and Computers, over here on the left, we see the name of my domain, CorpNet.xyz. As we pointed out just a second ago, that's the name of the domain that's being hosted on this domain controller.

If I expand my domain, I see a list of container objects. Container objects do just what their name implies: they contain other Active Directory objects. Their purpose is to keep things organized. Instead of putting all of our Active Directory objects within the domain, which would be really, really messy and confusing, we can sort the various objects into different containers to keep things nice, tidy, and well organized.

Notice when the name of a domain is selected over here, on the left, a list of all my container objects within the domain is displayed over here, on the right. Notice that there are several different types of container objects.

These first five container objects—Built-in, Computers, ForeignSecurityPrinciples, Manage Service Accounts, and Users—were all created by default when Active Directory was initially installed on this system. They were automatically populated with the various types of objects that they were designed to contain.

However, notice that there's another type of container object within the domain called an organizational unit. Within this domain, I have five different organizational units defined: Domain Controllers, MarketingUsers, ResearchUsers, SalesUsers, and TestUsers. With the exception of this container, Domain Controllers, these other organizational units were created by the system administrator.

The Domain Controller's organizational unit is a default container. It was created when Active Directory was initially deployed. It contains a computer account for this domain controller itself.

You might be asking, "What's the difference between an organizational unit and a container object?" Organizational units have a lot more flexibility; they can perform more functions. Container objects are very basic in nature. They're designed to hold the default objects that are created and used by Active Directory. The nice thing about these organizational units, as I've said before, is that they allow you to organize the various objects within your Active Directory domain.

There are many different approaches that you could take. One common approach is the one that I've used here, where I've defined an organizational unit for each of the divisions within my organization.

For example, I have an organizational unit for MarketingUsers. I have an organizational unit for ResearchUsers. I have an organizational unit for SalesUsers. I can put all the associated objects within Active Directory inside the appropriate container. For example, if I double-click ResearchUsers, it contains user accounts for all of the users in my company's research department.

This is one approach for creating Active Directory organizational units. But it's not the only one. Active Directory is very, very flexible. You could create different types of organizational units if you wanted.

For example, I've seen some organizations create organizational units based on geography. They might have one for North America, one unit for South America, and units for Europe, Africa, Australia, and Asia. You can configure the organizational units in whatever way works best for your organization.

For our purposes today, we're going to go with the approach that you see here, where we create a different organizational unit for each of the functional divisions within the organization. The users in Marketing, Research, and Sales are all set. Bust most organizations have more than three divisions.

If I need to create additional organizational units to support those divisions, I can click my domain, over here. Then I come over here, to this button, which allows us to create a new organizational unit. Click it, and let's create a new Organizational Unit for my Technical Support division within my company.

Note that there's an option to protect the container from accidental deletion. It's a really good idea to leave it on. It prevents the administrator from accidentally deleting the organizational unit. If you think about it, that would be really bad because the organizational unit typically contains a lot of user accounts. If you delete the organizational unit, you also automatically delete all of the user accounts within the organizational unit. If that's not what you were intending to do, that makes for a very bad day. I like to leave this option turned on. It keeps you from doing that.

Now my Technical Support organizational unit is defined. Within this organizational unit, I can create other Active Directory objects. I can create additional organizational units, meaning that I can create a hierarchical tree to reflect the structure of my organization.

For example, I can create a new organizational unit for the Help Desk group within my organization. I can also come back up here and create another organizational unit for my Backline Support group within the Technical Support division. Do you see how you can use these organizational units to mirror the structure of your company?
Within these organizational units, I can create additional Active Directory objects. The most common type of Active Directory object in an organizational unit are user accounts. These user accounts allow a given user to authenticate to the domain. They can be used to, say, log on to a workstation that's been joined to the domain.
You can also use these user accounts to control what the user is allowed to access. Basically, that user object is used to authenticate and to control access to domain resources. In other words, if a user does successfully log on to the domain, what level of access are they allowed to domain resources?
Let's suppose we need to create some user accounts for our Help Desk group. With Help Desk selected, I'll come over and click Create a New User in the current container. The first thing I need to do is provide a name. Roger is the first name, and Meinert is the last name. Then I need to provide a user logon name.
Usually, the user logon name you provide is dictated by your organization's security policy. For example, it may be something to the effect of first initial and last name. Click Next. Then we enter a password for that user. It's a bad security practice for the system administrator to know the user's password. For security purposes, only the user should know his or her password. Therefore, it's a good idea to turn this option on, User must change password at next logon. The first time the user logs on, they'll be prompted to change their password. Then you won't know what their password is.
Notice that you also have options to prevent a user from being able to change a password. You can also specify that the password never expires or that the account is disabled. These options can come in handy on occasion. Click Next. Click Finish.
Now we have a user object created in Active Directory. Let's go ahead and create a couple more. This first one will be Tom Czachor. Create the username. Click Next. Give Tom a password. Click Next and Finish. And just for practice, let's create one more: Dustin Johnson, username djohnson. Enter the password. Click Next and Finish one more time.
Now we have three user accounts for our Help Desk employees. We've talked about how to create organizational units. We've talked about how to create user accounts. You can also create group accounts using Active Directory Users and Computers. For example, we can create a group within the Help Desk organizational unit.
To do this, I'll click Create a New Group in the current container icon. Let's give it the name of Midnight Shift. This group will contain all the employees within the Help Desk that works on the midnight shift. Now we have a Midnight Shift group. To add these users to that group, we have a couple of different options. One of the easiest ways is to double-click the group, click Members, and click Add. Then enter in the users that we want to add.
First, we'll add Roger. Click OK. Click Add again. This time we'll add Tom. Type in his name and click OK. And we don't want to forget about Dustin. Let's make sure we add him. Click OK, and our three lucky midnight help desk employees are added. Click OK. Now Roger, Tom, and Dustin are all members of the Midnight Shift group.
Group accounts are really useful because, more than likely, all of the users within this group that work in the Help Desk group are going to need very similar levels of access to domain resources. Instead of going to each and every user account and making the same assignments over and over again on every user account, we can just make the assignments once, to the group, and then make all the users who need that level of access members of that group. This makes things really easy.
All right, that's all for our demo on managing Active Directory objects, specifically user objects.

4.4.4 Active Directory Facts

This lesson covers the following topics:

- Directory services
- Active Directory
- Active Directory components

Directory Services

A directory service stores information about users, computers, security groups/roles, and services. Each object in the directory has a number of attributes. The directory schema describes the types of attributes, what information they contain, and whether they are required or optional. In order for products from different vendors to be interoperable, most directory

services are based on the Lightweight Directory Access Protocol (LDAP) , which was developed from a standard called X.500.

Within an X.500-like directory, a distinguished name (DN) is a collection of attributes that define a unique identifier for any given resource. A distinguished name is made up of attribute-value pairs separated by commas. The most specific attribute is listed first, and successive attributes become progressively broader. This most specific attribute is the relative distinguished name, as it uniquely identifies the object within the context of successive (parent) attribute values.

A network directory lists the subjects (principally users, computers, and services), objects (such as directories and files) available on the network, and the permissions that subjects have over objects. A directory facilitates authentication and authorization, and it is critical that it be maintained as a highly secure service.

Active Directory

Active Directory is Microsoft's proprietary directory service. It is a centralized database that contains user accounts and security information. In a workgroup, security and management are decentralized. They occur on each computer, containing information about users and resources. With Active Directory, all computers share the same central database on a remote computer called a domain controller.

Active Directory is a hierarchical database. Hierarchical directory databases have the following advantages over a flat-file database structure:

Advantage	Description
Organization	Hierarchical databases let you sort and organize user accounts by location, function, and department.
Replication	The Active Directory database can be replicated to other systems. This eliminates the need to manually recreate user accounts on every system a user may need to access.
Delegation	Delegation allows you to assign users to manage portions of the Active Directory database without giving all users rights to the entire database. For example, you can assign an administrator to manage the sales department in North America and enable this administrator to create user accounts, remove user accounts, and change passwords. However, this sales administrator won't be allowed to access the accounting or development departments. As another example, you can enable an administrator to manage all departments in Europe but none in North America or Asia.
Scalability	A hierarchical database lets you grow the Active Directory to meet the needs of your environment.

Active Directory Components

Active Directory organizes network resources and simplifies management using the following components:

Component	Description
Domain	A domain is an administratively defined collection of network resources that share a common directory database and security policies. The domain is the basic administrative unit of an Active Directory structure. Depending on the network structure and requirements, the entire

Component	Description
	network might be represented by a single domain with millions of objects, or the network might require multiple domains.
Trees and forests	Multiple domains are grouped together in the following relationship: • A tree is a group of related domains that share the same contiguous DNS namespaces. • A forest is the highest level of the organization hierarchy and is a collection of related domain trees. The forest establishes the relationship between trees that have different DNS namespaces.
Organizational unit (OU)	An organizational unit is like a folder subdividing and organizing network resources within a domain. An organizational unit: • Is a container object. • Can hold other organizational units. • Can hold objects such as users and computers. • Can be used logically to organize network resources. • Simplifies security administration.
Generic container	Like OUs, generic containers organize Active Directory objects. Generic container objects: • Are created by default. • Cannot be moved, renamed, or deleted. • Have very few properties you can edit.
Object	Within Active Directory, each resource is identified as an object. Common objects include: • Users. • Groups. • Computers. • Shared folders. Each object contains additional information about the shared resource that can be used for locating and securing resources. Groups are composed of other directory objects with a common access level. The schema identifies the object classes (the type of objects) in the tree and the attributes (properties) of the objects. In Active Directory, each user is assigned a Security Account Manager (SAM) account name; therefore, each username must be unique.
Domain controller	A domain controller is a server that holds a copy of the Active Directory database that can be written to. Replication is copying changes to Active Directory between the domain controllers. In contrast, member servers are servers in the domain that do not have the Active Directory database.

4.4.5 Create OUs (Simulation)

Scenario

You are the IT administrator for a small corporate network. You have just installed Active Directory on a new Hyper-V guest server named CorpDC. Now, you need to create an Active Directory organizational unit (OU) structure based on the company's departmental structure.

In this lab, your task is to create the following organizational units (OUs) on the CorpDC server and ensure that each is protected from accidental deletion as follows:

- Beneath the CorpNet.local domain, create the following OUs:
  - Accounting
  - Admins
  - Marketing
  - Research-Dev
  - Servers
  - Support
  - Workstations
  - Sales
- Within the Sales OU, create the following OUs:
  - SalesManagers
  - TempSales

Explanation

While completing this lab, use the following information:

- Beneath the CorpNet.local domain, create the following OUs:
  - Accounting
  - Admins
  - Marketing
  - Research-Dev
  - Servers
  - Support
  - Workstations
  - Sales
- Beneath the Sales OU, create the following OUs:
  - SalesManagers
  - TempSales

Complete this lab as follows:

1. Access the CorpDC server.
   a. From the left pane of Hyper-V Manager, select **CORPSERVER** .
   b. From the Virtual Machines pane, double-click **CorpDC** .
2. Create the Active Directory organizational units (OUs) beneath the CorpNet.local domain.
   a. From the Server Manager's menu bar, select **Tools** > **Active Directory Users and Computers** .
   b. From the left pane, right-click **CorpNet.local** and then select **New** > **Organizational Unit** .

      You can also create OUs by selecting the **Create a new organizational unit in the current container** icon ( ) located in the Active Directory Users and Computers ribbon.

   c. Enter the ***name*** of the OU to be created.
   d. Ensure that **Protect container from accidental deletion** is selected and then select **OK** .
   e. Repeat steps 2b - 2d until all the required domain OUs are created.
3. Create the OUs within the Sales OU.

a. From the left pane, select **CorpNet.local** > **Sales** .
b. From the menu bar, select the **Create a new organizational unit in the current container** icon.
c. Enter the ***name*** of the OU to be created.
d. Ensure that **Protect container from accidental deletion** is selected and then select **OK** .
e. Repeat steps 3a - 3d to create the remaining OU.

4.4.6 Delete OUs (Simulation)

Scenario

You are the IT administrator for a corporate network. You have just installed Active Directory on a new Hyper-V guest server named CorpDC. You have created an Active Directory structure based on the company's departmental structure. While creating the structure, you added a Workstations OU in each of the departmental OUs. After further thought, you decide to use one Workstations OU for the entire company. As a result, you need to delete the departmental Workstations OUs.

In this lab, your task is to delete the following OUs on CorpDC:

- Within the **Marketing** OU, delete the **Workstations** OU.
- Within the **Research-Dev** OU, delete the **Workstations** OU.
- Within the **Sales** OU, delete the **Workstations** OU.

Explanation

To complete this lab, you need to delete the following OUs on CorpDC:

- Within the **Marketing** OU, delete the **Workstations** OU.
- Within the **Research-Dev** OU, delete the **Workstations** OU.
- Within the **Sales** OU, delete the **Workstations** OU.

Complete this lab as follows:

1. Access the CorpDC server.
   a. From Hyper-V Manager, select **CORPSERVER** .
   b. From the Virtual Machines pane, double-click **CorpDC** .
2. Delete the applicable OUs.
   a. From Server Manager, select **Tools** > **Active Directory Users and Computers** .
   b. Select **View** > **Advanced Features.**
   This enables the Advanced feature, allowing you to disable the OU from accidental deletion.
   c. From the left pane, expand **CorpNet.local** > ***the_parent OU*** .
   d. Right-click the ***OU*** that needs to be deleted and then select **Properties** .
   e. Select the **Object** tab.
   f. Clear **Protect object from accidental deletion** and then select **OK** .
   g. Right-click the ***OU*** to be deleted and then select **Delete** .
   h. Select **Yes** to confirm the OU's deletion.
   i. Repeat steps 2c - 2h to delete the remaining OUs.
3. From the *Active Directory Users and Computers* menu bar, select **View** > **Advanced Features** to turn off the Advanced Features view.

4.4.7 Group Policy (Lesson Video)

Transcript:

One of the great benefits of managing your networks with a domain is that a domain lets you apply various configuration settings such as security settings globally. You can use a Group Policy to configure settings across multiple Windows hosts in a domain.

In this lesson, we're going to spend some time looking at what Microsoft's Group Policy is and how it works.
Let's first explain just what a policy is.
A policy is a set of configuration settings that can be used to control the working environment of user accounts and computer accounts.
These settings can be applied in one of two ways.
First, an administrator could go to each computer and configure user and computer policies manually using the Local Group Policy Editor.
The disadvantage of using this method is that it takes a lot of time going from computer to computer and increases the possibility of user error.
Because of this, this method is typically only used for small networks that don't use Active Directory, or AD.
The better method for you to apply Group Polices is to configure and apply them using Active Directory's Group Policy Management.
Using this method, you're able to push down the needed configuration settings to the individual hosts within the domain straight from the domain controller.
As you can imagine, making the change once to a domain controller is a lot easier than going from computer to computer creating the same setting over and over.
It also means that you're less likely to make a mistake.
It's important to keep in mind that it's possible to use both local and domain policies at the same time.
But, if there's a conflict between the two, the domain policy will always take precedence over any local settings.
When using Group Policy, policies are grouped together in a collection and stored in a Group Policy object, or GPO.
Because GPOs contain a collection of settings, they're very powerful and you can use them to apply a number of things.
For example: registry settings, running scripts, applying templates, setting up software, making configuration settings, and more.
Here you see we're editing the GPO named Default Domain Policy, and notice that within this policy in the left panel we have a tree structure, in which we've stored our Group Policy settings.
In this example, if extend our tree out to the bottom, we go from: Computer Configuration to Policies; then to Windows Settings, to Security Settings, to Account Policies; and within Account Policies, we have a group of policies that are related to passwords.
If we click here within Password Policies, we have the group of policies as shown.
Notice that with these policies we can enforce such things as the password history.
Each policy will include some settings, and those can be seen here.
Double-clicking on one of these policies lets you configure these settings.
In the example seen here, Enforce password history is enabled and 24 passwords will be remembered. This means that the end user must use 24 unique passwords before they'll be allowed to reuse a password that they've used before.
As another example, you'll notice we also have a policy for setting the maximum password age. In other words, how many days the user can use the same password before they'll be forced to change it.
Right now, that's set to 42 days.
We also have additional policies to determine the password age, the password length, and whether passwords must comply with our complexity requirements.
As you can see, each of these individual policies lets us customize exactly how we want to manage our passwords.
Policies within a Group Policy object, such as the Default Domain Policy shown here, are divided into two different categories Computer Configuration and User Configuration.
You might be asking what the difference is between a Computer Configuration Policy and a User Configuration Policy.
A Computer Configuration Policy is applied to the computer itself. Those settings will be applied no matter who logs on to the computer. In other words, Group Policy doesn't care who the individual user is only which computer they're using.
The types of computer policies you may see under Computer Configuration include such things as: policies to control what software should be installed on the computer; scripts that might need to be run when the system boots or shuts down; password restrictions that are going to be applied to all the user accounts on the system; and also network security settings or registry settings that are going to be applied to the entire computer.

The important thing to remember about Computer Configuration Policies is that they're applied as the system boots up and initially connects to the domain. The system enforces them before any user logs on.
User Configuration Policies, on the other hand, are applied to a specific user. This means that a user's polices will still be applied regardless of which computer they use to log on to the domain.
As such, different policies can be set for different users. For example, you may have one computer system being used by three different users.
Taking advantage of user policies, you can apply different settings based upon the User Configuration Policy that you set up for individual users.
Some sample User Policy settings could include the software that's installed for a specific user, or scripts that run when they log on or log off.
It might also include security settings for your web browser. You could even use it to push down a list of favorites you've already selected for your browser.
You can also use it to push down registry settings that apply just to that individual user.
Just remember that the key difference between User Policies and Computer Policies is the that User Policies are applied on a per-user basis. So, they're not applied until the user logs on.
Computer Configuration Policies, on the other hand, are applied to the entire computer system no matter who logs on. So they're applied when the system initially boots up.
There are literally hundreds of different configuration settings that can be applied through Group Policy, and we can't possibly cover all of them here. But, there are a few key ones that you should be familiar with.
First, we have Account Policies. We use Account Policies to control such things as password settings, account lock out settings, Kerberos settings, and more.
Next, we have Local Policies/Audit Policy. You use the Audit Policy settings to configure auditing for various events such as log on, account management, or privileged use.
Next, we have Local Policies/User Rights assignment. User Rights determine what actions a given user can perform on a given computer.
We also have Local Policies/Security Policies. Security Policies are used to control such things as allowing a user to install an unsigned driver or requiring control-alt-delete to be pressed in order to log on.
We can also configure Registry Policies. You can use Registry Policies to configure specific registry keys and values. They can also specify whether you can read or even change a specific registry value.
Another set of Group Policies affect the file system. You can use File System Policies to configure file and folder permissions that apply to multiple computers.
We also have Software Restriction Policies. You can use Software Restriction Policies to define what software can run on any computer that's joined to the domain. These Software Restriction Policies can be applied to a specific user or they can be applied globally to all users.
Finally, we have administrative templates.
Administrative templates are registry-based settings that you can configure within a Group Policy object. They're very useful. You can forgo editing the registry or making configuration changes in Control Panel or the Settings app individually on every single computer in your network. Instead, just use the settings under administrative templates within Group Policy to control the computer configuration and centrally manage the user experience.
For example, you can use administrative templates to enable and configure various Windows features such as Bitlocker, offline files, or even parental controls.
You can also use administrative templates to do fun things like customize the start menu, customize the task bar, or even set up the desktop environment.
That's it for this lesson.
In this video, we talked about Group Policy management.
We first talked about what Group Policy is and how policies can be used to push down configuration settings to all the computers in your domain.
We talked about the relationship between Group Policy and Local Policy settings that you can configure on individual machines.
We then looked at the difference between Computer Configuration Policies and User Configuration Policies.
And then we ended this lesson by reviewing a list of policies that you should be familiar with in order to effectively manage a domain.

4.4.8 Use Group Policy (Demo Video)

Transcript:

In this demonstration, we're going to spend some time working with Group Policy. Group Policy contains a myriad of settings that you can use to control the way Windows looks, runs, and behaves.

Understand that there are two different sets of Group Policy settings for Windows. If your Windows system isn't a part of a domain, the locally stored Group Policy settings are applied. But if your system is a member of a domain, you can also apply Group Policy settings from a domain controller.

Let's begin by looking at the first option. To start, I'm going to come down here and click on Search. Then I'm going to type in gpedit.msc and press Enter to run the Group Policy Editor.

You'll notice up in the window's title bar that it says Local Group Policy Editor. That word Local is key. This tells me that the settings I'm configuring are just for this computer system. You'll also notice that there are two kinds of settings. We have Computer Configuration, and we have User Configuration.

A Computer Configuration Group Policy is applied when the system first boots up, and it's applied to every single user who logs in. The User Configuration settings are different. They're applied on a per-user basis, and they're applied only when the user logs in.

Let's take a look at some sample Computer Configuration settings. We'll go to Administrative Templates > Windows Components and then to Store. Within the Store set of Group Policy settings, you can see each individual setting here. For example, there's one here that turns off the Store Application. I'll double-click it. When I do, I can look down here in Help, and it tells me what this setting does. Basically, this either allows or denies access to the Windows Store application. If this system is being used in a company environment, you might not want your users going out and downloading whichever apps they want from the Windows Store. In this situation, you could turn Store access off by clicking Enabled and OK.

Now the policy is enabled, but my Store icon is still available on the taskbar. I could still go out to the Windows Store because this change won't be applied until the system shuts down and reboots again. That's because it's a computer policy. Okay, let's go ahead and close this.

Now let's go down here under User Configuration and expand Administrative Templates. Then let's click Control Panel. There are many different Group Policy settings we can configure here that affect the way the user interacts with Control Panel.

Remember that the user settings aren't applied until the user actually logs in. For example, we could come over here and say, "You know what, I don't want my users having any access to Control Panel or to my Settings app." Your average user could really mess things up if they don't know what they're doing.

So what we want to do is go in and turn access completely off. We're going to enable this policy. It tells us down here that, once enabled, if a user tries to select a Control Panel item, a message will appear explaining that this action isn't allowed. The next time somebody logs in to this system, they'll no longer have access to Control Panel.

So that's one way that you can apply Group Policy settings—with the Local Group Policy Editor. This works fine if you're on a network of three, four, or five Windows systems. But you'd still have to go from system to system to make all the changes. If you're in a big network with thousands of computers, this isn't a good way to do things.

A better way is to configure your Group Policy settings on a domain controller. Using a domain controller, you can make configuration changes like we just did here, but all at once for everyone. They'll automatically be pushed down to each individual Windows system that's a member of the domain. We can't make those changes here from a workstation, though. Instead, we have to switch over and go to the domain controller server itself. Let's do that.

Alright, I'm at my DC1 domain controller. I can come over here to Tools and click Group Policy Management. When I do, I can browse down through my domain to the Group Policy Object that I want to manage. I right-click and click Edit. This is the Default Domain Policy. It's configured by default as soon as I set up Active Directory, and it'll be applied by default to all Windows systems that are members of the domain unless I tell it not to.

You'll notice that Domain Group Policy is divided up basically the same way as the Local Group Policy. You have your Computer Policies and your User Policies. It works in the exact same way, too—Computer Policies are applied as soon as the system comes online, and User Policies are applied when the user logs in.

Again, the nice thing is that the changes I make on the domain controller will be applied to all the Windows systems that are members of that domain. For example, let's go into my Computer Configuration > Windows Settings > Security Settings > Account Policies and then to Account Lockout Policy.

The Account Lockout Policy is used to prevent an attacker from trying to guess passwords on the system. It does this by saying, "You know what, you're allowed to enter the wrong password a limited number of times. Once you exceed that number, I'm going to lock the system because the authorized user is probably not going to take that many attempts to log in to their own system."

And let's come down here to Account lockout threshold. You'll notice that by default, Account Lockout Policy is disabled. You can try as many passwords as you want, and the system won't lock you out. That's not a good thing. It allows an attacker to just sit there at the console and try password after password until they get the right one.
So let's say that if you can't get the right password after three attempts, we're going to assume that you're not the right person and lock the system. Click OK. When I do, it says, "You know what, if we're going to do that, we need to change the values of these other Group Policy settings as well." These are things like how long we're going to keep the system locked and how long we're going to wait before we reset the counter that's recording the number of invalid login attempts.
So now I'm going to give the user 3 chances to log in, and it's going to stay locked for 30 minutes. After 30 minutes, we'll reset the lockout counter. And because it's a Computer Configuration, the next time each domain system powers on and connects, the setting will be automatically pushed down to them. I don't have to walk around from system to system. Making changes just once is really convenient.
Okay, that's it for this demonstration. In this demo, we configured Group Policy settings.

4.4.9 Group Policy Facts

This lesson covers the following topics:

- Group Policy
- GPO structure

Group Policy

A policy is a set of configuration settings applied to users or computers. Group policies allow the administrator to apply multiple settings to multiple objects within the Active Directory domain at one time. A set or collection of Group Policy configurations is called a Group Policy Object (GPO). The GPO is a collection of files that includes registry settings, scripts, templates, and software-specific configuration values.

GPO Structure

Each GPO has a common structure with hundreds of configuration settings that can be enabled and configured. Settings are divided into two categories:

GPO Category	Description
Computer Configuration	Computer policies are enforced for the entire computer and are initially applied when the computer boots. Computer policies are in effect regardless of the user logging into the computer. Computer policies include: • Software that should be installed on a specific computer • Scripts that should run at startup or shutdown • Password restrictions that must be met for all user accounts • Network communication security settings • Registry settings that apply to the computer (the HKEY_LOCAL_MACHINE subtree) Computer policies also include a special category of policies called user rights. User rights identify system maintenance tasks and the users or groups who can perform these actions. Actions include: • Changing the system time

	• Loading and unloading device drivers • Removing a computer from a docking station • Shutting down the system Computer policies are initially applied as the computer boots and are enforced before any user logs on.
User Configuration	User policies are enforced for specific users and are applied when the user logs on. User Policy settings include: • Software that should be installed for a specific user • Scripts that should run at logon or logoff • Internet Explorer user settings (such as favorites and security settings) • Registry settings that apply to the current user (the HKEY_CURRENT_USER subtree) User policies are initially applied as the user logs on and customizes Windows based on his or her preferences.

GPOs apply to objects when they are linked to containers and configured with specific settings.

- GPOs can be linked to Active Directory domains or organizational units (OUs).
- Built-in containers (such as the Computers container) and folders cannot have GPOs linked to them.
- A GPO applied to an OU affects only the users and computers in that OU.
- A GPO applied to a domain affects all users and computers in all OUs in that domain.
- A local GPO is stored on a local machine. It can be used to define settings even if the computer is not connected to a network.
- A specific setting in a GPO can be:
    - Undefined - this means that the GPO has no value for that setting and does not change the current setting.
    - Defined - this means that the GPO identifies a value to enforce.
- GPOs are applied in the following order:

1. The Local Group Policy on the computer.
2. GPOs linked to the site.
3. GPOs linked to the domain that contains the User or Computer object.
4. GPOs linked to the organizational unit(s) that contain(s) the User or Computer object (from the highest-level OU to the lowest-level OU).

Use the acronym LSDOU (Local, Site, Domain, and OU) to help you remember the order that GPOs are applied. The local policy is always applied, but it may be overwritten by a policy from Active Directory.

- Individual settings within all GPOs are combined to form the effective Group Policy setting as follows:
    - If a setting is defined in one GPO and undefined in another, the defined setting is enforced (regardless of the position of the GPO in the application order).
    - If a setting is configured in two GPOs, the setting in the last-applied GPO is used.

4.4.10 Create and Link a GPO (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You would like to use Group Policy to enforce settings for certain workstations on your network. You have prepared and tested a security template file that contains policies that meet your company's requirements.

In this lab, your task is to perform the following on CorpDC:

- Create a GPO named **Workstation Settings** in the CorpNet.local domain.
- Link the Workstation Settings GPO to the following organizational units (OUs):
    - **Marketing** > **TempMarketing**
    - **Sales** > **TempSales**
    - **Support**
- Import the **ws_sec.inf** template file located in C:\Templates to the Workstation Settings Group Policy object.

Explanation

While completing this lab, use the following information:

- Link the Workstation Settings GPO to the following organizational units (OUs):
    - **Marketing** > **TempMarketing**
    - **Sales** > **TempSales**
    - **Support**
- Import the **ws_sec.inf** template file located in **C:\Templates** .

Complete this lab as follows:

1. Access the CorpNet.local domain.
    a. From Server Manager, select **Tools** > **Group Policy Management** .
    b. Expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** .
    c. Maximize the window for better viewing.
2. Create the Workstation Settings GPO and link it to the CorpNet.local domain.
    a. Right-click the **Group Policy Objects** OU and select **New** .
    b. In the Name field, use **Workstation Settings** and then select **OK** .
3. Link OUs to the Workstation Settings GPO.
    a. Right-click the ***OU*** and select **Link an Existing GPO** .
    b. Under Group Policy Objects, select **Workstation Settings** from the list and then select **OK** .
    c. Repeat steps 3a-3b to link the additional OUs.
4. Import the ws_sec.inf security policy template.
    a. Expand **Group Policy Objects** .
    b. Right-click **Workstation Settings** and select **Edit** .
    c. Under Computer Configuration, expand **Policies** > **Windows Settings** .
    d. Right-click **Security Settings** and select **Import Policy** .
    e. Browse to **C:\Templates** .
    f. Select **ws_sec.inf** and then select **Open** .

4.4.11 Create User Accounts (Simulation)

Scenario

You are the IT administrator for a small corporate network. You recently added an Active Directory domain to the CorpDC server to manage network resources centrally. You now need to add user accounts in the domain.

In this lab, your task is to create the following user accounts on CorpDC:

User	Job Role	Departmental OU
Juan Suarez	Marketing manager	Marketing\MarketingManagers
Susan Smith	Permanent sales employee	Sales\PermSales
Mark Burnes	Sales manager	Sales\SalesManagers
Borey Chan	Temporary sales employee	Sales\TempSales

Use the following user account naming standards and specifications as you create each account:

- Create the user account in the departmental OU corresponding to the employee's job role.
- User account name: ***First name*** + ***Last name***
- Logon name: ***firstinitial*** + ***lastname*** with **@CorpNet.local** as the domain
- Original password: **asdf1234$** (must change after the first logon)
- Configure the following for the temporary sales employee:
    - Limit the logon hours to allow logon only from **8:00 a.m.** to **5:00 p.m.** , **Monday** through **Friday** .
    - Set the user account to expire on **December 31st** of the current year.

Explanation

Use the following user account specifications as you create each account.

User	Job Role	Departmental OU
Juan Suarez	Marketing manager	Marketing\MarketingManagers
Susan Smith	Permanent sales employee	Sales\PermSales
Mark Burnes	Sales manager	Sales\SalesManagers
Borey Chan	Temporary sales employee	Sales\TempSales

Complete this lab as follows:

1. Access Active Directory Users and Computers on the CorpDC server.
    a. From Hyper-V Manager, select **CORPSERVER** .
    b. From the Virtual Machines pane, double-click **CorpDC** .
    c. From the Server Manager's menu bar, select **Tools** > **Active Directory Users and Computers** .
    d. Maximize the window for better viewing.

2. Create the domain user accounts.
    a. From the left pane, expand **CorpNet.local** .
    b. Browse to the appropriate ***OU*** .
    c. Right-click the ***OU*** and select **New** > **User** .
    d. In the *First name* field, enter the user's ***first name*** .
    e. In the *Last name* field, enter the user's ***last name*** .
    f. In the User logon name field, enter the user's ***logon name*** , which should be the first letter of the user's first name together with their last name *(e.g., jsuarez).*

        The domain @CorpNet.local is appended automatically to the end of the logon name.

    g. Select **Next** .
    h. In the *Password field* , enter **asdf1234$** .
    i. In the *Confirm password* field, enter **asdf1234$** .
    j. Ensure **User must change password at next logon** is selected and then select **Next** .
    k. Select **Finish** to create the object.
    l. Repeat steps 2b–2k to create the additional users.
3. Modify user account restrictions for the temporary sales employee.
    a. Right-click **Borey Chan** and select **Properties** .
    b. Select the **Account** tab.
    c. Select **Logon hours** .
    d. From the Logon Hours dialog, select **Logon Denied** to clear the allowed logon hours.
    e. Select the time range of **8:00 a.m.** to **5:00 p.m.** , **Monday** through **Friday** .
    f. Select **Logon Permitted** to allow logon.
    g. Select **OK** .
    h. Under Account expires, select **End of** .
    i. In the *End of* field, use the drop-down calendar to select **31 December** of the ***current year*** .
    j. Select **OK** .

4.4.12 Manage User Accounts (Simulation)

Scenario

You are the IT administrator for a small corporate network. You recently added an Active Directory domain on the CorpDC server to manage network resources centrally. Organizational units in the domain represent departments. User and computer accounts are in their respective departmental OUs.

Over the past few days, several personnel changes have occurred that require changes to user accounts.

In this lab, your task is to use the following information to make the necessary user account changes on CorpDC:

- Mary Barnes from the Accounting Department has forgotten her password, and now her account is locked.
    - Unlock the account.
    - Reset the password to **asdf1234$**
    - Require a password change at the next logon.
- Mark Woods has been fired from the accounting department. Disable his account.
- Pat Benton is returning to the Research-Dev department from maternity leave. Her account is disabled to prevent logon. Enable her account.
- Andrea Simmons from the Research-Dev department has recently married.
    - Rename the account **Andrea Socko** .
    - Change the last name to **Socko** .
    - Change the display name to **Andrea Socko** .
    - Change the user logon and the pre-Windows 2000 user logon name to **asocko** .

- For all users in the Support OU (but not the SupportManagers OU), allow logon only to the Support computer.

Explanation

Complete this lab as follows:

1. Access Active Directory Users and Computers on the CorpDC server.
    a. From the Server Manager's menu bar, select **Tools** > **Active Directory Users and Computers** .
    b. Maximize the window for better viewing.
2. Unlock the Mary Barnes account.
    a. From the left pane, expand and select **CorpNet.local** > **Accounting** .
    b. Right-click **Mary Barnes** and select **Reset Password** .
    c. In the New password field, enter **asdf1234$** .
    d. In the Confirm password field, enter **asdf1234$** .
    e. Make sure **User must change password at next logon** is selected.
    f. Make sure **Unlock the user's account** is selected.
    g. Select **OK** .
    h. Select **OK** to confirm the change.
3. Disable the Mark Woods account.
    a. Right-click **Mark Woods** and select **Disable Account** .
    b. Select **OK** to confirm the change.
4. Enable Pat Benton's account.
    a. From the left pane, select **Research-Dev** .
    b. From the right pane, right-click **Pat Benton** and select **Enable Account** .
    c. Select **OK** to confirm the change.
5. Rename the Andrea Simmons account.
    a. Right-click **Andrea Simmons** and select **Rename** .
    b. Enter **Andrea Socko** and press **Enter** . This opens the Rename User dialog.
    c. In the Last name field, enter **Socko** .
    d. In the User logon name field, replace the old name with **asocko** .
    e. Select **OK** .
6. Configure user account restrictions.
    a. From the left pane, select **Support** .
    b. Press the **Ctrl** key, and then from the right pane, select both the **Janice Rons** and **Tom Plask** users to edit multiple users at the same time.

        In Safari, press **Command** and select each user.

    c. Right-click the ***user accounts*** and select **Properties** .
    d. Select the **Account** tab.
    e. Select **Computer restrictions** .
    f. Select **Log On To** .
    g. Select **The following computers** .
    h. In the Computer name field, type **Support** .
    i. Select **Add** .
    j. Select **OK** .
    k. Select **OK** .

4.4.13 Create a Group (Simulation)

Scenario

You are the IT administrator for the CorpNet domain. You have decided to use groups to simplify the administration of access control lists. Specifically, you want to create a group containing the department managers.

In this lab, your task is to use Active Directory Users and Computers to complete the following actions on the CorpDC server:

- In the Users container, create a group named **Managers** . Configure the group as follows:
  - Group scope: **Global**
  - Group type: **Security**
- Make the following users members of the Managers group:
-

Organization Unit	Username
Accounting	**Mark Woods**
Research-Dev	**Pat Benton**
Marketing\MarketingManagers	**Juan Suarez**
Research-Dev\ResearchManagers	**Arlene Kimbly**
Sales\SalesManagers	**Mark Burnes**
Support\SupportManagers	**Shelly Emery**

Explanation

Complete this lab as follows:

1. Access Active Directory Users and Computers on the CorpDC server.
   a. From the Server Manager's menu bar, select **Tools** > **Active Directory Users and Computers** .
   b. Maximize the window for better viewing.
2. In the Users container, create a group named **Managers** .
   a. From the left pane, expand and select **CorpNet.local** > **Users** .
   b. Right-click the **Users** container and select **New** > **Group** .

      You can also create a new group by selecting the Create a new group in the current container icon found in the ribbon.

   c. In the *Group name* field, enter **Managers** .
      A pre-Windows 2000 group name is created automatically but can be changed.
   d. Under Group scope, make sure **Global** is selected.
   e. Under Group type, make sure **Security** is selected and select **OK** .
3. Add user accounts to the Managers group.
   a. From the left pane, ensure that the Users container is still selected.
   b. From the right pane, right-click **Managers** and select **Properties** .

c. Select the **Members** tab.
d. Select **Add** .
e. In the *Enter the object names to select* field, enter all the ***usernames* .** Use a semicolon to separate each name.
   Example: Steve Hoffer; Peter Williams; Princess Diana
f. Select **Check Names** .
g. Select **OK** to add the users and close the dialog.
h. Select **OK** to close the Managers Properties dialog.

You can also add individual users to a group by right-clicking the user and selecting **Add to a group** .

4.4.14 Create Global Groups (Simulation)

Scenario

You are the IT Administrator for the CorpNet.local domain. You are in the process of implementing a group strategy for your network. You have decided to create global groups as shadow groups for specific departments in your organization. Each global group will contain all users in the corresponding department.

In this lab, your task is to:

- Create the following global security groups on the CorpDC server in their corresponding OUs:

OU Creation Location	New Group Name
Accounting	Accounting
Research-Dev	Research-Dev
Sales	Sales

- Add all user accounts in the corresponding OUs and sub-OUs as members of the newly created groups.

Explanation

While completing this lab, use the following information:

OU Creation Location	New Group Name
Accounting	Accounting
Research-Dev	Research-Dev

Sales	Sales

Complete this lab as follows:

1. Access Active Directory Users and Computers on the CorpDC server.
    a. From the Server Manager's menu bar, select **Tools** > **Active Directory Users and Computers** .
    b. Maximize the window for better viewing.
    c. From the left pane, expand **CorpNet.local** .
2. Create the groups.
    a. Right-click the **OU** where the new group is to be added and select **New** > **Group** .
    b. In the *Group name* field, enter the **name** of the group.
    c. Make sure the **Global** *Group scope* is selected.
    d. Make sure the **Security** *Group type* is selected.
    e. Select **OK** .
3. Add users to groups.
    a. In the right pane, right-click the ***user account(s)*** and select **Add to a group** . (Use the Ctrl or Shift keys to select and add multiple user accounts to a group at one time.)
    b. In the *Enter the object names to select* field, enter the **name** of the group.
    c. Select **Check Names** and verify that the object name was found.
    d. Select **OK** to accept the groups added.
    e. Select **OK** to acknowledge the change.
    f. If a sub-OU with users exists, double-click on the sub-OU and then repeat step 3. Do this for each sub-group.
4. Repeat steps 2 - 3 for additional groups and users.

4.4.15 Practice Questions (Section Quiz)

q_actdir_ad_01_secp8

What is the name of the service included with the Windows Server operating system that manages a centralized database containing user account and security information?

Answers:

- Active Directory
- Active directory
- active directory
- AD
- ad

Explanation:

Active Directory (AD) is a centralized database that is included with the Windows Server operating system. Active Directory is used to store information about a network. It stores such things as user accounts, computers, printers, and security policies.

q_actdir_ad_02_secp8

Match each Active Directory term on the left with its corresponding definition on the right.

Answers:

- Tree
- Forest
- Domain
- Organizational unit
- Object

Explanation:

The Active Directory structure includes the following components:

- A tree is a group of related domains that share the same contiguous DNS namespace.
- A forest is a collection of related domain trees.
- A domain is an administratively defined collection of network resources that share security policies and a common directory database.
- An organizational unit (OU) is like a folder. An OU subdivides and organizes network resources within a domain.
- An object is a network resource as identified within Active Directory.

q_actdir_ad_advantages_secp8

Which of the following are advantages of using hierarchical databases like Active Directory? (Select two.)

Answers:

- They allow for decentralized security and management.
- ***They enable replication of the database to other systems.**
- They limit the growth of the Active Directory to meet the needs of your environment.
- ***They allow for organization of user accounts by location, function, and department.**
- They require manual recreation of user accounts on every system a user may need to access.

Explanation:

The following are advantages of using hierarchal databases like Active Directory:

- They enable replication of the database to other systems. Explanation: hierarchical databases like Active Directory can replicate the database to other systems, eliminating the need to manually recreate user accounts on every system a user may need to access.
- They allow for organization of user accounts by location, function, and department. Explanation: hierarchical databases let you sort and organize user accounts by location, function, and department, providing a structured and organized approach to managing user accounts.

The following are not advantages of using hierarchal databases like Active Directory:

- They allow for decentralized security and management. One of the key advantages of hierarchical databases like Active Directory is that they provide centralized security and management, not decentralized.
- They limit the growth of the Active Directory to meet the needs of your environment. Hierarchical databases like Active Directory are scalable, meaning they allow for growth to meet the needs of your environment, not limit it.
- They require manual recreation of user accounts on every system a user may need to access. One of the key advantages of hierarchical databases like Active Directory is that they can replicate the database to other systems, eliminating the need for manual recreation of user accounts.

q_actdir_centralized_database_secp8

What is the primary function of Active Directory as a centralized database in a network?

Answers:

- ***It stores and organizes all user accounts and security information.**
- It provides internet access to all computers in the network.
- It serves as a backup system for all files in the network.
- It manages the power supply to all computers in the network.

Explanation:

Active Directory is Microsoft's proprietary directory service. It is a centralized database that contains user accounts and security information. In a workgroup, security and management are decentralized. They occur on each computer, containing information about users and resources. With Active Directory, all computers share the same central database on a remote computer called a domain controller.

While Active Directory can manage certain network services, providing internet access is not its primary function. This is typically handled by network routers and switches.

Active Directory does not serve as a backup system for all files in the network. Its primary function is to manage user accounts and security information. Backup systems are separate and are used to store copies of data that can be restored in the event of data loss.

Active Directory does not manage the power supply to all computers in the network. Its primary function is to manage user accounts and security information. Power management is typically handled by each individual computer's operating system and hardware.

q_actdir_domain_controller_secp8

Which of the following BEST describes the domain controller component of Active Directory?

Answers:

- ***A domain controller is a server that holds a copy of the Active Directory database that can be written to and is responsible for copying changes to Active Directory between the domain controllers.**
- A domain controller is a user account that has administrative privileges to manage the Active Directory database.
- A domain controller is a specific type of network resource within a domain.
- A domain controller is a software application that manages the replication of the Active Directory database.
- A domain controller is a physical device that connects the network to the Active Directory database.

Explanation:

A domain controller is a server that holds a writable copy of the Active Directory database. It is responsible for managing changes to the database and replicating these changes to other domain controllers to ensure consistency across the network.

A domain controller is not a user account, but a server that manages the Active Directory database.

A domain controller is not a type of network resource, but a server that manages the Active Directory database.

While software is involved in the process, a domain controller is not a software application but a server that holds a copy of the Active Directory database.

A domain controller is not just a physical device that connects the network to the Active Directory database, but a server that holds a writable copy of the Active Directory database and manages changes to it.

q_actdir_domain_grouping_secp8

A large multinational corporation has multiple domains that share the same contiguous DNS namespaces, as well as domains with different DNS namespaces. The IT department is tasked with organizing these domains.

Which of the following options best describes how the domains should be grouped?

Answers:

- All domains should be grouped into a single tree, regardless of their DNS namespaces.
- ***Domains with the same contiguous DNS namespaces should be grouped into a tree, and all trees should be grouped into a forest.**
- Domains with different DNS namespaces should be grouped into a tree, and all trees should be grouped into a forest.
- Domains with the same contiguous DNS namespaces should be grouped into a forest, and all forests should be grouped into a tree.
- All domains should be grouped into a single forest, regardless of their DNS namespaces.

Explanation:

Domains with the same contiguous DNS namespaces should be grouped into a tree, and all trees should be grouped into a forest. In Active Directory, a tree is a group of related domains that share the same contiguous DNS namespaces. A forest, on the other hand, is the highest level of the organization hierarchy and is a collection of related domain trees. The forest establishes the relationship between trees that have different DNS namespaces.

The following are incorrect answers:

- All domains should be grouped into a single tree, regardless of their DNS namespaces. This is incorrect because a tree in Active Directory is a group of related domains that share the same contiguous DNS namespaces. Domains with different DNS namespaces should not be grouped into the same tree.
- Domains with different DNS namespaces should be grouped into a tree, and all trees should be grouped into a forest. This is incorrect because a tree in Active Directory is a group of related domains that share the same contiguous DNS namespaces. Domains with different DNS namespaces should not be grouped into the same tree.
- Domains with the same contiguous DNS namespaces should be grouped into a forest, and all forests should be grouped into a tree. This is incorrect because a forest in Active Directory is a collection of related domain trees, not the other way around.
- All domains should be grouped into a single forest, regardless of their DNS namespaces. This is incorrect because while a forest can contain multiple trees with different DNS namespaces, it is not necessary or always beneficial to group all domains into a single forest. The organization of domains into trees and forests should be based on the specific needs and structure of the organization.

q_actdir_generic_containers_secp8

Which of the following statements correctly describe the characteristics of generic containers in Active Directory? (Select two.)

Answers:

- Generic containers can be moved, renamed, or deleted.
- ***Generic containers are created by default.**

- Generic containers have numerous properties you can edit.
- Generic containers cannot hold other organizational units.
- ***Generic containers are used to organize Active Directory objects.**

Explanation:

The following are statements that correctly describe generic containers:

- Generic containers are created by default in Active Directory.
- Generic containers are used to organize Active Directory objects. Like organizational units, generic containers are used to organize Active Directory objects.

The following are incorrect statements describing generic containers:

- Generic containers can be moved, renamed, or deleted.
- Generic containers have numerous properties you can edit. Generic containers have very few properties you can edit.
- Generic containers cannot hold other organizational units. While it's true that generic containers are different from organizational units, the statement is misleading because it suggests that generic containers cannot contain other objects, which is not true. Generic containers can contain other Active Directory objects, just like organizational units.

q_actdir_hierarchal_secp8

What is one of the main advantages of Active Directory being a hierarchical database?

Answers:

- It allows for faster internet speeds.
- ***It allows for organization and sorting of user accounts and resources.**
- It allows for automatic software updates.
- It allows for increased storage capacity.

Explanation:

One of the main advantages of Active Directory being a hierarchical database is that it allows for the organization and sorting of user accounts and resources. This can be done by location, function, department, or any other criteria that suits the needs of the organization.

The hierarchical nature of Active Directory does not directly impact internet speeds. Internet speeds are typically determined by the network infrastructure and internet service provider.

While Active Directory can be used to manage software updates, this is not a direct result of it being a hierarchical database. Software updates are typically managed through specific services or tools, such as Windows Server Update Services (WSUS).

The hierarchical nature of Active Directory does not increase storage capacity. Storage capacity is determined by the physical storage hardware and the file system in use.

q_actdir_schema_secp8

A new IT administrator is tasked with managing Active Directory for their company. The administrator needs to understand the types of objects in the tree and the properties of these objects.

Which of the following BEST describes the Active Directory component that the new administrator needs to understand?

Answers:

- The administrator needs to understand the domain controller, as it manages the Active Directory database.
- ***The administrator needs to understand the schema, as it identifies the object classes and their attributes in the tree.**
- The administrator needs to understand the forest, as it is the highest level of the organization hierarchy.
- The administrator needs to understand the organizational unit, as it subdivides and organizes network resources within a domain.
- The administrator needs to understand the replication process, as it copies changes to Active Directory between the domain controllers.

Explanation:

The schema in Active Directory identifies the object classes (the types of objects) in the tree and the attributes (properties) of these objects. This is what the administrator needs to understand to manage the types of objects and their properties in Active Directory.

While understanding the domain controller is important, it does not directly help the administrator understand the types of objects and their properties in the tree.

While understanding the forest is important, it does not directly help the administrator understand the types of objects and their properties in the tree.

While understanding the organizational unit is important, it does not directly help the administrator understand the types of objects and their properties in the tree.

While understanding the replication process is important, it does not directly help the administrator understand the types of objects and their properties in the tree.

q_gpo_computer_01_secp8

You want to ensure that all users in the Development OU have a common set of network communication security settings applied.

Which action should you take?

Answers:

- Create a GPO user policy for the Development OU.
- ***Create a GPO computer policy for the computers in the Development OU.**
- Create a GPO folder policy for the folders containing the files.
- Create a GPO computer policy for the Computers container.

Explanation:

Network communication security settings are configured in the Computer Policies section of a GPO.

Built-in containers (such as the Computers container) and folders cannot be linked to a GPO.

q_gpo_computer_02_secp8

You have several computers running Windows 11. The computers are members of a domain.

For all computers, you want to remove access to administrative tools from the Start menu and hide notifications from the system tray.

What should you do?

Answers:

- ***Use Group Policy**
- Use account restrictions
- Use account policies
- Use file screens

Explanation:

Use Group Policy to control the desktop for groups of users or computers. For example, you can prevent access to specific desktop or Start menu features.

Account policies are specific Group Policy settings that control user passwords. Account restrictions are settings applied in the user account that restrict logon hours or computers.

You use file screens to control the types of files that can be saved within a folder.

q_gpo_linked_01_secp8

A user has complained about not being able to remove a program that is no longer needed on a computer. The Programs and Features page is not available in Control Panel.

You suspect that a policy is enabled that hides this page from the user. But after opening the Local Group Policy Editor, you see that the Hide Programs and Features page is set to Not configured. You know that other users in this domain can access the Programs and Features page.

To determine whether the policy is enabled, where should you look next?

Answers:

- ***GPOs linked to organizational units that contain this user's object.**
- GPOs linked to the domain that contains this user's object.
- The Default Domain Policy GPO.
- The Local Group Policy.

Explanation:

You should look at GPOs linked to organizational units that contain this user's object to see where the Hide Programs and Features page policy might be enabled.

If the policy was enabled in a GPO linked to the domain, it would be applied to all users in the domain.

The next level GPOs are applied from is the level of GPOs linked to organizational units that contain the user's object.

q_gpo_linked_02_secp8

The Hide Programs and Features page setting is configured for a specific user as follows:

Policy	Setting
Local Group Policy	Enabled
Default Domain Policy GPO	Not configured
GPO linked to the user's organizational unit	Disabled

After logging in, the user is able to see the Programs and Features page. Why does this happen?

Answers:

- ***The GPO linked to the user's organizational unit is applied last, so this setting takes precedence.**
- The Default Domain GPO is applied last. It is set to Not configured, so it doesn't change the configuration.
- The Local Group Policy is applied last. It is set to Enabled, which makes the Programs and Features page visible.
- The GPO linked to the user's organizational unit is applied first, so this setting takes precedence over settings that are applied later.

Explanation:

The GPO linked to the user's organizational unit is applied last. With this in mind, the setting that disables the policy to hide the Programs and Features page takes precedence.

In this scenario, Local Group Policy enables the policy to hide the Programs and Features page.

When the Default Domain Policy GPO is applied, this policy is set to Not configured. It doesn't change anything.

When the GPO linked to the user's organizational unit is applied, the setting for this policy is disabled. This reverses the setting in the Local Group Policy and makes the Programs and Features page visible to the user.

The Local Group Policy is applied first. GPOs linked to the user's domain are applied second and take precedence over settings in the Local Group Policy. GPOs linked to the user's organizational unit are applied last and take precedence over any preceding policy settings.

q_gpo_lsdou_secp8

Which of the following is the correct acronym to remember the order in which Group Policy Objects (GPOs) are applied?

Answers:

- OSDL
- ***LSDOU**
- DOLS
- SLOD

Explanation:

LSDOU (Local, Site, Domain, Organizational Unit) is correct. The order of application is Local, Site, Domain, and Organizational Unit.

OSDL (Organizational Unit, Site, Domain, Local) is incorrect. The order of application starts with Local, not Organizational Unit.

DOLS (Domain, Organizational Unit, Local, Site) is incorrect. The order of application starts with Local, not Domain.

SLOD (Site, Local, Organizational Unit, Domain) is incorrect. The order of application starts with Local, not Site.

q_gpo_order_01_secp8

Which statement is true regarding the application of GPO settings?

Answers:

- If a setting is defined in the Local Group Policy on the computer and not defined in the GPO linked to the OU, the setting is not applied.
- If a setting is not defined in the Local Group Policy and is defined in the GPO linked to the OU, the setting is not applied.
- If a setting is defined in the Local Group Policy on the computer and defined differently in the GPO linked to the OU, the Local Group Policy setting is applied.
- ***If a setting is defined in the Local Group Policy on the computer and not defined in the GPO linked to the OU, the setting is applied.**

Explanation:

GPOs are applied in the following order:

1. The Local Group Policy on the computer.
2. GPOs linked to the domain that contains the User or Computer object.
3. GPOs linked to the organizational unit(s) that contain(s) the User or Computer objects (from the highest-level OU to the lowest-level OU).

Individual settings within all GPOs are combined to form the effective Group Policy setting as follows:

- If a setting is defined in one GPO and undefined in another, the defined setting is enforced (regardless of the position of the GPO in the application order).
- If a setting is configured in two GPOs, the setting in the last-applied GPO is used.

q_gpo_order_02_secp8

Group Policy Objects (GPOs) are applied in which of the following orders?

Answers:

- ***Local Group Policy, GPO linked to site, GPO linked to domain, GPO linked to organizational unit (highest to lowest).**
- GPO linked to site, GPO linked to domain, GPO linked to organizational unit (highest to lowest), Local Group Policy.

- Local Group Policy, GPO linked to site, GPO linked to domain, GPO linked to organizational unit (lowest to highest).
- GPO linked to site, GPO linked to domain, GPO linked to organizational unit (lowest to highest), Local Group Policy.

Explanation:

Group Policy Objects (GPOs) are applied in the following order:

- The Local Group Policy on the computer.
- GPOs linked to the site.
- GPOs linked to the domain that contains the User or Computer object.
- GPOs linked to the organizational unit (OU) that contains the User or Computer object (from the highest-level OU to the lowest-level OU).

q_gpo_user_secp8

You manage an Active Directory domain. All users in the domain have a standard set of internet options configured by a GPO linked to the domain, but you want users in the Administrators OU to have a different set of internet options.

What should you do?

Answers:

- Create a GPO computer policy for the Administrators OU.
- ***Create a GPO user policy for the Administrators OU.**
- Create a Local Group Policy on the computers used by members of the Administrators OU.
- Create a GPO user policy for the domain.

Explanation:

Internet options are configured in the User Policies section of a GPO. Linking this policy to the Administrators OU only applies it to users in that OU because GPOs linked to OUs are applied last.

If Local Group Policies are created on the Administrator's computers, the policies are overwritten by the GPO that is linked to the domain, which applies a standard set of internet options to all users in the domain. There is already a GPO user policy linked to the domain.

4.5 Hardening Authentication

As you study this section, answer the following questions:

- What does the minimum password age setting prevent?
- What is a drawback to account lockout for failed password attempts?
- What are the advantages of a self-service password reset management system?

In this section, you will learn to:

- Configure user account restrictions
- Configure account policies and UAC settings
- Use password managers

- Configure account password policies
- Hardening user accounts
- Restrict local accounts
- Secure default accounts
- Enforce user account control
- Configure smart card authentication

The key terms for this section include:

Term	Definition
Multifactor authentication	Using more than one method to authenticate users.
Smart cards	Similar in appearance to credit cards, smart cards have an embedded memory chip that contains encrypted authentication information. These cards are used for authentication.
Microprobing	The process of accessing a smart card's chip surface directly to observe, manipulate, and interfere with the circuit.
Radio frequency identification (RFID)	The wireless, non-contact use of radio frequency waves to transfer data.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts.
	• Authentication, Authorization, and Accounting (AAA)
	  ○ Authenticating people
	2.4 Given a scenario, analyze indicators of malicious activity.
	• Indicators
	  ○ Account lockout
	  ○ Concurrent session usage
	  ○ Blocked content
	  ○ Impossible travel
	  ○ Resource inaccessibility
	4.5 Given a scenario, modify enterprise capabilities to enhance security.
	• Operating system security
	  ○ Group Policy
	4.6 Given a scenario, implement and maintain identity and access management.
	• Provisioning/de-provisioning user accounts

Exam	Objective

	• Permission assignments and implications
	• Multifactor authentication
	  ○ Implementations
	  ▪ Security keys
	  ○ Factors
	  ▪ Something you have
	• Password concepts
	  ○ Password best practices
	  ▪ Length
	  ▪ Complexity
	  ▪ Reuse
	  ▪ Expiration
	  ▪ Age
	  ○ Password managers
	  ○ Passwordless
	• Privileged access management tools
	  ○ Just-in-time permissions
	  ○ Password vaulting
	  ○ Temporal accounts

4.5.1 Hardening Authentication (Lesson Video)

Transcript:

In this lesson, we're going to talk about hardening authentication. Hardening means strengthening. You want to make sure your authentication methods are strong so you can be confident that the users accessing your network are who they claim to be.

As simple as it may sound, user education is one of the most important parts of authentication hardening. A surprisingly large percentage of security issues within an organization can be tracked back to an employee or another authorized user. While some of these breaches are intentional, many are unintentional; they're a result of a seemingly harmless action.
After all, the strongest doors and the toughest locks become inconsequential if an employee inadvertently opens the door for a potential attacker because they're trying to be helpful.
End user security training should be extensive. For the purpose of authentication hardening, encourage your employees to use password creation methods that they'll remember. Remind them that they shouldn't write down their passwords. Teach them how to recognize social engineering attempts, and remind them not to share their physical or digital access with anyone for any reason.
Next, let's look at passwords more closely. Since passwords are a primary method of authorization, be sure to harden your password policies.
For example, configuring password aging policies causes passwords to be valid for only a specified length of time. After this time has passed, the user will need to change their password.
Enforcing password history prevents the end user from being able to reuse the same password over and over. This can be set for a certain number of times. You might configure your settings so employees can't reuse any of their past five passwords. Or you could set up an indefinite policy so that a user can never use any of the same passwords they've used in the past.
Another way to harden passwords is to enforce password complexity. Doing this prevents users from creating simple passwords like "1234" or "password." When you configure your enforcement policy, you can require capital letters, numbers, or special characters. You can even require that the password be a minimum number of characters long.

You can also allow end users to manage their own passwords without requiring help from the system administrator. This is known as self-service password management. While this can be a convenient option for the user and the administrator, be sure to consider whether it's suitable for your desired level of security.
As we've already discussed, passwords are a common authentication method, and they do provide a layer of security. But just because your network or data is protected by a passwords doesn't mean it's secure.
Whenever possible, you should use multifactor authentication. This means using more than one method to authenticate your users.
End users can authenticate in many ways.
The user could provide something they know, like a password or a PIN.
They could provide something they have, like an ID card or a security token. They could provide something that they are by using their biometric data, such as a palm print, a fingerprint, or a retina scan. Robust authentication processes use two or more of these factors.
Some administrators mistakenly assume that requiring users to supply a username and a password constitutes multifactor authentication. That's incorrect. A username and a password are both items that the end user knows, so they constitute single-factor authentication. If you want to make your authentication process more robust, you can add an additional authentication factor.
For example, you could require a name and password along with a smart card: something that the user knows and something that the user has.
Or you could require a fingerprint scan and a PIN: something that the user is and something that the user knows.
Another way to harden your authentication is to apply restrictions to user accounts. First, you can limit the number of concurrent logins to one per person. This way, if someone forgets to log off of a device, they can't log into a second device without terminating the first connection.
You can also limit when an end user can log in to the system. For example, you may specify that end users can log in between the hours of seven in the morning and six at night. If they try to log in outside of that window, they're blocked.
Another thing you can do is set an account lockout threshold. The account lockout threshold determines the number of failed login attempts that can occur before the user account is locked. Or you could configure a reset lockout counter. The reset lockout counter parameter specifies the amount of time that must pass after a failed login attempt before the failed login counter is reset. This is to slow down password guessing as well as dictionary or brute force attacks.
While we're on the topic of user accounts, let's discuss monitoring and maintenance. Most network administrators are sitting on a pile of data about their users.
There are logs to show when users are signing in, when they're entering incorrect passwords, when they're resetting passwords, where they log in from, and more.
Keeping an eye on activity logs can be kind of dull and time-consuming. So instead of doing it manually, a lot of people use software to do it for them. If unusual or risky behavior is detected, the software will notify you so can act accordingly. This could mean requiring an additional method of authentication, or it could mean locking an account completely.
Maintaining accounts is also important. Some large corporation networks have been breached using outdated or inactive user accounts.
If an employee leaves the organization, that user's account needs to be disabled or deleted. You should also disable inactive accounts. These are commonly associated with contract employees, who only access the network periodically. If the user doesn't need access for a time, disable their account. When the user needs access again, you can re-enable it.
Automatic account expiration is another option where the account is only enabled for a certain amount of time. After that time has passed, the account is automatically disabled. This is especially useful for temporary employees.
You should also make sure that the user accounts have the appropriate level of access. On a Windows network, there are two separate sets of user accounts that you must manage. The first set of accounts resides on the domain controller, and they're managed by the admins.
The second type of accounts are the local user accounts on each individual workstation. These are local accounts. You may not want users to have local administrative access. If they do, they could adjust local file system permissions, install malware, or even create a backdoor user account. Because of this the default administrator account should be the only administrator-level local user accounts on your workstation, and you may want to rename it to obscure it. All other users should be standard users that aren't members of the Administrators group.
You should also disable the local guest user account on each workstation. This account allows someone to authenticate and gain limited access to the system without providing any type of credentials.
Now, let's talk about remote access. The most important thing you can do to harden your remote access is to ensure that it's only enabled for end users who actually need it. Remote-access clients should connect to the internal network through a demilitarize zone, or DMZ. This allows traffic monitoring, which lets you verify that remote access connections are authorized and legitimate. You can also restrict remote access to only certain authorized IP addresses.

In addition, you can limit concurrent logins, such as only one per user. Every time a remote user connects, their last login date and time should be displayed. This allows the end user to see if somebody else has been using their connection to gain access to the network.
Lastly, you should audit remote logins. Check to see if remote connections are made at suspicious times or from suspicious IP addresses, and check for a spike in the number of failed logins. All these events indicate that some type of attack is probably being conducted.
And that's it for this lesson. In this video, we discussed hardening authentication through user education, stronger passwords, multifactor authentication, account restrictions, account monitoring, account maintenance, and limited remote access.

4.5.2 Configure User Account Restrictions (Demo Video)

Transcript:

In this demonstration, we'll look at hardening a Windows workstation by configuring user account restrictions. By setting these restrictions in a domain, they will always be applied, even if the user logs in from a different system, or if a different user logs into system.

Let's begin by opening Active Directory Users and Computers. Let's drill down into the Employees OU so we can see our users.
First, we need to do several things to harden these accounts. Let's choose the temporary employee, Daniel Jackson.
There are several things we can do to harden his account, specifically because he's a temporary employee.
This user works Monday through Friday from 8:00 in the morning until 5:00 at night. We need to restrict the user's logon hours and days to the times he is working. Also, because the user is a temporary employee, we should configure the account to expire automatically when the employment term expires.
We want to disable the account automatically in case we forget to manually disable it.
First, let's go into Account settings > Logon Hours. Everything that is highlighted in blue is a permitted logon block or hour.
You can see the days here, Sunday through Saturday, and anything highlighted in white will be logon denied. In this case, we want to limit logon to Monday through Friday, 8:00 a.m. through 5:00 p.m. If we select Saturday and select Logon Denied, it denies logon for the whole day. It whites it all out. We'll select Sunday, as well.
We'll select all the days from midnight 8:00 a.m and disable those. Then we'll go from here to 5:00 p.m to midnight and disable those hours. Now he only has login permitted from 8 a.m. to 5 p.m. Monday through Friday. We click OK, and that applies the restriction.
We also need to set the account expiration. You can see it's set to Never. We want to change this. Let's say this account will expire at the end of the year. We'll click December 31. We'll click Apply. Click OK.
Now Daniel Jackson's account will automatically deny him access unless he's logging in Monday through Friday, 8:00 a.m. to 5:00 p.m. The account will be disabled automatically on December 31, at the end of the day.
Also, in this OU, we have an employee who's on extended leave.
Samantha Carter has been out on maternity leave, so her account is disabled. We did not delete her account because we are expecting her to come back to work, and we don't want to have to recreate all her permissions when she returns.
Samantha is returning tomorrow, so let's re-enable her account.
To do this, come up here and right-click the account, and enable account. Notice that the little black arrow on the user icon disappeared – that means the account is now enabled again.
The last thing we want to do in this demonstration is restrict a user to a particular computer on the network. The user will be allowed to log on only from that system. Let's say the user, Samantha Carter, is allowed to log in only on the lab computer.
Let's go to Properties. We'll go to Account and click Log On To... As you can see here, the user can currently log in to all computers on the domain. We don't want that because the user should only be in the lab.
We'll set it manually. Here, you type in the computer name after selecting the following computers. You type in the computer's net bios, domain name, or DNS name.
Let's type in 'LabComputer' and click Add. Notice this is the only computer listed now for the user. So, this is the only computer the user can log in to.
You can add more than one computer. Let's click OK. Now this user can log in only to that lab computer.

That's it for this demonstration. In this demonstration, we showed you how to configure user account restrictions. We first looked at how to restrict login times for a given user account.
We disabled and enabled user accounts. We set an expiration date for an account. Then we ended the demonstration by configuring a user account to use a specific computer in the network.

4.5.3 Configure Account Policies and UAC Settings (Demo Video)

Transcript:

In this demonstration, we'll show you actions you can take to harden Windows authentication. First, we'll configure account policies that will be applied to all the users in the domain.

Then, we'll enforce UAC settings. An easy way to enforce UAC settings is to make the configuration changes to a GPO in the domain. When you use a GPO, settings are automatically enforced on each workstation whenever a user logs in to the domain.
Using Group Policy, you make the changes once, instead of having to do it manually on each workstation.
Let's go to Tools > Group Policy Management and right-click Default Domain Policy. Click Edit. That opens the Group Policy Management Editor. Then, let's go to Computer Configuration > Policies > Window Settings > Security Settings > Account Policies > Password Policy.
Here, you see the password policies that are available.
First, let's change the minimum password length from seven characters to a more secure 10 characters.
The value you choose for many of these settings depends on your environment, the organization, and how you want to balance security with ease of use for the users.
If you set the password policy too high, it might be difficult for users to manage. We'll set the password length at 10 because that's a reasonable number of characters for a password.
Now, let's go to Enforce Password History. The value is set to 24. This means the system keeps a 24-password history. That prevents users from repeating recent passwords.
This policy makes sure that users can't reuse a password until they have made 24 different password changes. Click OK. That's fine where it's set.
Let's change our maximum password age. We're going to change that to 30 days.
A minimum password age of one day means the user could change the password every day. Some users might try changing the password every day, so they get back to an old password after 24 days of passwords changes. We don't want to allow that.
Let's change the minimum password age to 10 days. That means users have to keep a password for 10 days.
Now, let's look at password complexity requirements. These are basic complexity requirements that Windows has defined.
In most cases when you click the Explain tab, it'll provide more information about the policy setting. Here you can see all the password requirements that Windows has set. So, we'll click OK.
After configuring password policies, let's configure the account lockout policies. If we go here, the first one we'll look at is the account lockout threshold.
You can see that it's currently set to 0 invalid logon attempts. Let's double-click it. Zero invalid logon attempts mean it's not enabled. It's defined, but a policy set at zero is, by definition, not enabled.
Let's change this to 3. That means, after three invalid password entries, the account will lock. Click OK.
If you look, this automatically sets the other two values for the account lockout policy to 10 minutes each. Let's go into these.
The Account Lockout duration determines the length of time before a locked account will automatically be unlocked. The Reset Account Lockout Counter determines how long an invalid password attempt will be counted toward the Account Lockout Threshold.
10 minutes is fine for both of these settings, so I'm going to leave them as-is.
The next thing we'll look at is the User Account Control. If we go to Local Policies > Security Options, we can look at the UAC settings.
A word of caution – configuring the UAC incorrectly can cause some real problems. It can generate excessive elevation prompts, it can cause applications not to install correctly, and be a source of frustration for your users and your help desk. For this video, we're going to go through each of these settings and configure the group policy according to how the UAC behaves by default.

It's always a good idea to put new settings into a new group policy object. So, let's close this and go back to the Group Policy Management. We'll go into the Group Policy Objects folder and right click, then select New to create a new GPO. We'll call this UAC Settings.
We'll edit this GPO and go to Computer Configuration, Policies, Windows Settings, Security Settings, Local Policies, Security Settings, and scroll down to the bottom.
You can read more about what each of these settings do in the description. Here is how we are going to configure them for this video.
Admin Approval Mode for the Built-in Administrator account will be enabled.
Allow UIAccess applications to prompt for elevation without using the Secure Desktop will be disabled.
Behavior of the elevation prompt for administrators in Admin Approval Mode will be set to "Prompt for consent for non-Windows binaries".
Behavior of the elevation prompt for standard users will be set to "Prompt for credentials on the secure desktop".
Detect application installations and prompt for elevation will be enabled.
Only elevate executables that are signed and validated will be disabled.
Only elevate UIAccess applications that are installed in secure locations will be enabled.
Run all administrators in Admin Approval Mode will be enabled.
Switch to the secure desktop when prompting for elevation will be enabled.
And last, virtualize file and registry write failures to per-user locations will be enabled.
Now that I have the GPO configured the way I want, I just need to apply it to an OU, and all the computers beneath that OU will inherit the settings from it.
That concludes this demo. In this demonstration, we hardened Windows authentication.

4.5.4 Use Password Managers (Demo Video)

Transcript:

With the large number of passwords you may accumulate from websites, computer logins, or even network equipment, you may consider using a password manager to store your passwords securely. Most password managers offer a way to configure two-factor authentication, so in the event a hacker obtains your password for the password manager, they'd be stopped by the two-factor authentication process used to log in to the password manager.

Today, we're going to look at a product called Bitwarden. This is a very popular password manager that offers excellent protection for your passwords. Best of all, the starting plan for this is free. There are some options, as you can see here, that offer better protection if you decide to pay. Some of these options include Advanced 2FA, which is the same as two-factor authentication. If you have a business, you can use a paid business option to store your business passwords.
We've already created our Bitwarden account; however, when you create your account, two-factor authentication (2FA) isn't set up yet. To set this up, we must go to our account options and select Account settings. On the Security tab, there will be an option for Two-step login. The paid options allow the use of Yubikey, Duo, or WebAuthn. For this account, we'll set up an authenticator app and use Authy. It will first prompt for the master password just to make sure you want to do this.
This next screen is important. From here, you'd download your authenticator app on your phone. Two examples are Authy and Google Authenticator. Although I can't show the app part, once installed, you can bring up the app, add an account, and then use your phone's camera to scan the QR code to simply add Bitwarden to Authy. Once added, use the latest six-digit code from Authy to finish this setup.
Now that we have 2FA set up let's test it out. If we Log out of our account and then re-enter our email address and password, you'll see that it will now prompt for a 2FA code instead of just logging you in. After we input our six-digit code, we're successfully logged in.
Our authentication is configured properly, so let's add some passwords. All we have to do is click New Item. We can define what this item is, like a Login, Card, Identity, or a Secure note. For now, let's stick with Login. Next, we need to give it a name, configure a username, and input the password. Keep in mind this will save the password for whatever login you've already configured, say from Outlook, Yahoo, or even your favorite forum site.
In order for a browser extension to pick up the password, you must set a URL. We'll see why in the next step. If, by chance, you wish to write a note about this password, it gives you that option. Let's Save our password.
This next part makes it super easy to log in to websites. Most password managers have a browser extension, so let's add that. Going to the three dots and selecting Extensions allows us to find the one we're looking for. After typing

Bitwarden, this is the one we want. Extensions don't always show up on the toolbar, so let's change that setting. Going back to our Extension settings, we can choose to make Bitwarden show in the toolbar. This makes it a lot easier to access passwords. When the extension is set up, we'll be required to re-login with our username, password, and code from our authenticator app. Regardless of whether it's logging into Bitwarden on a website, extension, or even a phone app, it should be the same process.
The main thing we want to see is if the username and password get populated in the fields presented on a website, so let's test this. Two ways we can get to the site are to manually type it in or use the arrow button to launch the URL from the password. Either way, you'll get to the same place. Once you click Sign in, you'll be prompted for your email. You can see that when we go up to our Bitwarden extension, we have one available login. When clicked, it will populate the fields on the screen. It will also do the same for the password—granted, this isn't a real login; however, you get the idea of how it can auto-populate a username and password without typing it in each time.
Let's go back to Bitwarden to show some other features. Say you'd like to put a secure note in; what we could do is add a New Item. In this case, we're going to select a Secure note. Say this is my combination lock for my shed, and I don't want to forget this combination. We could write a note for this and then Save it. Now, you can see how useful this password manager can be in securing all your information.
That's it for this demo. In this demo, we showed you how to set up and configure a password manager.

4.5.5 Configure Account Password Policies (Simulation)

Scenario

You have been asked to perform administrative tasks for a computer that is not a member of a domain. To increase security and prevent unauthorized access to the computer, you need to configure specific password and account lockout policies.

In this lab, your task is to use the Local Security Policy to configure the following password and account lockout policies:

- Configure password settings so that the user must:
    - Cycle through ten passwords before reusing an old one.
    - Change the password every 90 days.
    - Keep the password for at least 14 days.
    - Create a password at least eight characters long.
    - Create a password that meets complexity requirements, such as using uppercase letters, lowercase letters, numbers, or symbols.
- Configure the account lockout policy to:
    - Lock out any user who enters five incorrect passwords.
    - Unlock an account automatically after 60 minutes.
    - Configure the number of minutes that must elapse after a failed logon attempt to 10 minutes.
-

Explanation

In this lab, your task is to edit the Local Security Policy and configure settings as follows:

Policy Location	Policy	Setting
Account Policies/Password Policy	**Enforce password history**	**10**
	Maximum password age	**90**

	Minimum password age	**14**
	Minimum password length	**8**
	Passwords must meet complexity requirements	**Enabled**
Account Policies/Account Lockout Policy	**Account lockout threshold**	**5**
	Account lockout duration	**60**
	Reset account lockout counter after	**10**

Complete this lab as follows:

1. Using Windows Administrative Tools, access the Local Security Policy.
   a. Select **Start** .
   b. Locate and expand **Windows Administrative Tools** .
   c. Select **Local Security Policy** .
   d. Maximize the window for better viewing.
2. Configure the password policies.
   a. From the left pane, expand and select **Account Policies** > **Password Policy** .
   b. From the center pane, expand the **Policy** column.
   c. Double-click the ***policy*** to be configured.
   d. Configure the ***policy settings*** .
   e. Select **OK** .
   f. Repeat steps 2c-2e to configure the additional password policies.
3. Configure the account lockout policies.
   a. From the left pane, select **Account Lockout Policy** .
   b. From the center pane, expand the **Policy** column.
   c. Double-click the ***policy*** to be configured.
   d. Configure the ***policy settings*** (if needed, answer any prompts shown).
   e. Select **OK** .
   f. Repeat steps 3c-3e to configure the additional lockout policies.

4.5.6 Hardening User Accounts (Demo Video)

Transcript:

In this demonstration, we'll look at things you can do to harden Windows workstation authentication within your network. First, we'll look at restricting local accounts.

We want to ensure that only authorized users have administrative rights to the local Windows operating system. If a workstation is joined to a domain, then we typically use a domain user account to authenticate.
However, we sometimes forget that there are still local user accounts on each individual Windows workstation. One or more of them may have administrative access to the local systems; this presents a security issue.
A better option is to configure these settings in a GPO in the domain. This will allow you to make these configuration settings once and have them automatically enforced on every workstation.

To begin, we'll go to Tools > Group Policy Management > Group Policy Objects, and create a new GPO. We'll call this "Restrict Admin Group Membership". Go to Edit.
As you know, when you install a Windows workstation, several user accounts are created by default. You define administrator, guest, and user accounts during the installation process.
We want to harden authentication by renaming the default administrator account to something that makes it less obvious that the account is an administrator.
We also want to disable the guest account. It provides a degree of access to the system without a password. We also want to look for any local user account that has the Password never expires option enabled.
We'll do this with our group policy. We'll navigate to Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > Security Options. We're going to enable the option to rename the Administrator account, we'll rename the administrator account to 'Elvis'. And we'll set the Guest Account Status option so this GPO will disable the Guest account.
Next, we need to make sure that administrative rights are only given to the groups we define. We're going to use a group policy feature called Restricted Groups. Use caution when adding entries in Restricted Groups. We're going to create an entry that defines who is in the Administrators group. When you define a group membership with a Restricted Group setting, it will evict any member of the group that isn't defined by your group policy. This is a highly effective way to make sure admin rights do not get assigned outside of group policy. However, there is a risk you will evict an account that needs admin rights. If that happens, you just need to be prepared for it, and define that account with the group policy setting we're about to show you.
We'll navigate to Computer Configuration > Policies > Security Settings > Restricted Groups. We'll right click and create a new entry for "BUILTIN\Administrators".
We also want to add our domain admins to this group.
And we'll add one more entry for the "Elvis" account – that we renamed our Administrator account.
Let's click OK. Now we can see this is created right here. Next, we want to secure the default accounts on each workstation.
All user accounts should have their passwords changed on a regular schedule.
We also want to look for any user accounts that have never been used. The default for creating a new local user account on a Windows workstation should have the option User must change password at next logon enabled.
Therefore, we're going to assume that any user account that has that option marked has never been used because the user has never logged in. We want to get rid of all accounts that have never been used.
To do this, we'll switch over to the Windows 11 workstation. We'll right-click in the window and go to Computer Management.
Now we're going to go to Local Users and Groups > Users. This shows the users currently on the system.
The next thing we want to look at are accounts that have never been logged in to. If we double-click this account right here, for Mr. Hank Hill, you can see the option to require password change at login is not checked. That's good.
This means that this user has logged in before. If we look at this one, now you can see this says User must change password at next login.
This means this user has never logged in, which means this account is just sitting here and, possibly, is an account that's not needed. Most likely, that's the case. We'll delete this account because we probably don't need it.
The other thing you want to look for is the Password never expires checkmark. We'll check our TestOut account.
We'll leave this in for our purposes. But normally, you don't want a normal user to have a password that never expires.
If you do have this option enabled on your Windows workstations, you probably want to uncheck it and require that the user has to change the password, just like normal accounts on the domain. For now, we're going to click OK.
That's it for this demonstration. In this demo, we hardened Windows authentication for Windows workstations in a network.

4.5.7 Restrict Local Accounts (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You are working to increase the authentication security of the domain. You need to ensure that only authorized users have administrative rights to all local machines. Local users and groups can be controlled through a GPO linked to the domain.

In this lab, your task is to edit the Default Domain Policy and configure the Local Users and Groups policy settings as follows:

- Create a policy to update the built-in Administrator local group.
- Delete all member users.
- Delete all member groups.
- Add **BUILTIN\Administrator** to the group.
- Add **%DOMAINNAME%\Domain Admins** to the group.

The policy you create should remove all members of the built-in Administrators group and then add only the members specified. Use **BUILTIN\Administrator** and **%DOMAINNAME%\Domain Admins** in the policy to indicate which accounts to add.

Explanation

Complete this lab as follows:

1. Access the **CorpNet.local** domain under Group Policy Management.
    a. From Server Manager, select **Tools** > **Group Policy Management** .
    b. Maximize the windows for better viewing.
    c. Expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** .
2. Create a policy to update the built-in Administrator local group.

a. Right-click **Default Domain Policy** and select **Edit** .
    b. Maximize the windows for better viewing.
    c. Under Computer Configuration, expand **Preferences** > **Control Panel Settings** .
    d. Right-click **Local Users and Groups** and select **New** > **Local Group** .
    e. Using the *Group name* drop-down, select **Administrators (built-in)** .
    f. Select **Delete all member users** to remove all member users.
    g. Select **Delete all member groups** to remove all member groups.
    h. Select **Add** .
    i. In the *Name* field, enter **BUILTIN\Administrator** and then select **OK** .
    j. Select **Add** .
    k. In the *Name* field, enter **%DOMAINNAME%\Domain Admins** and then select **OK** .
    l. Select **OK** to save the policy.

4.5.8 Secure Default Accounts (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. You are improving office computers' security by renaming and disabling default accounts.

In this lab, your task is to perform the following on the Office1 computer:

- Rename the Administrator account, **Yoda** .
- Disable the Guest account.
- Verify that *Password never expires* is not selected for any local users. This forces them to change their passwords regularly.
- Delete any user accounts with *User must change password at next logon* selected. This indicates that a user has never logged in.

Explanation

Complete this lab as follows:

1. Access the computer's Computer Management tool.
    a. Right-click **Start** and select **Computer Management** .
    b. Under System Tools, expand **Local Users and Groups** .
    c. Select **Users** .
2. Rename the Administrator account.
    a. From the center pane, right-click **Administrator** and select **Rename** .
    b. Enter **Yoda** and press **Enter** .
3. Disable the Guest account.
    a. Right-click **Guest** and select **Properties** .
    b. Select **Account is disabled** and select **OK** .
4. Remove Password never expires option if it is selected.
    a. Right-click a ***user*** and select **Properties** .
    b. Deselect **Password never expires** (if selected) and then select **OK** .
    c. Make a note of any user who has *User must change password at next logon* .
    d. Repeat steps 4a-4c for each user.
5. Delete any unused accounts.
    a. Right-click the ***user*** that has **User must change password at next logon** selected and select **Delete** .
    b. Select **Yes** to confirm the deletion of the account.

4.5.9 Enforce User Account Control (Simulation)

Scenario

You are the IT administrator for a small corporate network. The company has a single Active Directory domain named CorpNet.local. You need to increase the domain's authentication security. You need to make sure that User Account Control (UAC) settings are consistent throughout the domain and in accordance with industry recommendations.

In this lab, your task is to configure the following UAC settings in the Default Domain Policy on CorpDC as follows:

User Account Control	Setting
Admin Approval mode for the built-in Administrator account	**Enabled**
Allow UIAccess applications to prompt for elevation without using the secure desktop	**Disabled**
Behavior of the elevation prompt for administrators in Admin Approval mode	**Prompt for credentials**
Behavior of the elevation prompt for standard users	**Automatically deny elevation requests**
Detect application installations and prompt for elevation	**Enabled**
Only elevate UIAccess applications that are installed in secure locations	**Enabled**
Only elevate executables that are signed and validated	**Disabled**
Run all administrators in Admin Approval mode	**Enabled**

Switch to the secure desktop when prompting for elevation	**Enabled**
Virtualize file and registry write failures to per-user locations	**Enabled**

User Account Control policies are set in a GPO linked to the domain. In this scenario, edit the Default Domain Policy and configure settings in the following path:
Computer Configuration > **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** .

Explanation

While completing this lab, use the following information when configuring the UAC settings.

User Account Control	Setting
Admin Approval mode for the built-in Administrator account	**Enabled**
Allow UIAccess applications to prompt for elevation without using the secure desktop	**Disabled**
Behavior of the elevation prompt for administrators in Admin Approval mode	**Prompt for credentials**
Behavior of the elevation prompt for standard users	**Automatically deny elevation requests**
Detect application installations and prompt for elevation	**Enabled**
Only elevate executables that are signed and validated	**Disabled**
Only elevate UIAccess applications that are installed in secure locations	**Enabled**
Run all administrators in Admin Approval mode	**Enabled**
Switch to the secure desktop when prompting for elevation	**Enabled**
Virtualize file and registry write failures to per-user locations	**Enabled**

Complete this lab as follows:

1. On CorpDC, access the **CorpNet.local** domain for Group Policy Management.
    a. From Server Manager, select **Tools** > **Group Policy Management** .
    b. Maximize the window for easy viewing.
    c. Expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** .
2. Configure the UAC settings.
    a. Right-click **Default Domain Policy** and select **Edit** .
    b. Maximize the window for better viewing.
    c. Under Computer Configuration, expand and select **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** .
    d. From the right pane, double-click the ***policy*** you want to edit.

e. Select **Define this policy setting** .
f. Select **Enable** or **Disable** as necessary.
g. Edit the ***value*** for the policy as needed, and then select **OK** .
h. Repeat steps 2d–2g for each policy setting.

4.5.10 Hardening Authentication Facts

This lesson covers the following topics:

- Hardening authentication methods
- Hardening authentication best practices

Hardening Authentication Methods

Hardening means to strengthen. You want to make sure your authentication methods are strong so that you can be confident that users accessing your network are who they say they are.

The following table provides various methods for strengthening your authentication.

Method	Description
Password Policies	Account policies help you control the composition and use of passwords. Password policies include: • Enforce password history - This determines the number of unique new passwords that have to be used before an old password can be reused. This helps to prevent users from reusing any recent passwords. • Maximum password age - This requires users to change their password after a given number of days. • Minimum password age - This determines the number of days that a password must be used before the user can change it. This prevents users from reverting back to their original password immediately after they have changed it. • Minimum password length - This identifies the minimum number of characters in a password. • Password must meet complexity requirements - A complex password prevents using passwords that are easy to guess or crack. Complex passwords must meet the following minimum requirements: ○ Cannot contain the user's account name or parts of the user's full name that exceed two consecutive characters ○ Must be at least six characters in length ○ Must contain characters from three of the following four categories: ▪ English uppercase characters (A through Z) ▪ English lowercase characters (a through z) ▪ Base-10 digits (0 through 9) ▪ Non-alphabetic characters (for example, !, $, #, or %) Complexity requirements are enforced when passwords are changed or created.

Multifactor Authentication	When possible, multifactor authentication should be used. This means using more than one method to authenticate your users. End users can be authenticated using three types of factors: • Something you know • Something you have • Something you are Robust authentication processes use two or more of these factors.
Account Restrictions	Account restrictions place restrictions on the use of a user account for login. For example, you can: • Prohibit multiple concurrent logins • Allow logins only during certain days and hours • Allow logins only from specific computers • Create expiration dates for user accounts for temporary users to prevent them from being used past a certain date
Account Monitoring	Account monitoring can help you detect unusual or risky behavior. You should monitor for the following: • Login activity. • Suspicious logins for the user (spikes, logins at unusual time of day, and/or frequent or failed logins). • Remote-access traffic.
Account Maintenance	The following list provides best practices for account maintenance: • Delete an employee's account when the employee leaves the organization. • Disable inactive accounts. • Use automatic account expiration when applicable. • Restrict remote access only to authorized clients (filtering by IP address).
Limit Remote Access	The following precautions should be taken when administering remote access: • Allow remote access to the network only for those users who need it to perform their duties (not standard for all users). • Do not allow remote access clients to connect directly to the internal network. Allow remote access clients to connect to a DMZ and then monitor the traffic. • Restrict remote access only to authorized clients . You can filter by IP address.
Account Lockout Policies	Account lockout disables a user account after a specified number of incorrect login attempts. Account lockout policies include:

	• Account lockout duration - Specifies the number of minutes a locked-out account remains locked out before automatically becoming unlocked. When set to 0, an administrator must unlock the account. • Account lockout threshold - Specifies the number of failed logon attempts that causes a user account to be locked out. • Reset account lockout counter after - Specifies the number of minutes that must elapse after a failed logon attempt before the failed logon attempt counter is reset to 0 bad logon attempts. For example, if this value is set to 60 minutes and the account lockout threshold is set to 5, the user can enter up to four incorrect passwords within one hour without the account being locked. Account lockout can be used to prevent attackers from guessing passwords, but it can also be used maliciously to lock an account and prevent a valid user from logging in.

Hardening Authentication Best Practices

When controlling user account and password security, be aware of the following:

- For large environments, implement a password management system with a self-service password reset management system. This allows a user to change his or her own password and ensures that only he or she knows it. In a system where administrators hand out passwords that users cannot change, passwords lack security. In this type of arrangement, no matter how complex the password is, more than one person knows what it is. This can affect the security of the system.
- Implement account auditing to track incorrect login attempts. Small numbers of incorrect logon attempts occur naturally as users mistype or forget passwords. Large numbers of incorrect login attempts could identify a potential hacker trying to guess passwords.
- Scan systems to identify unused user accounts or accounts with blank passwords.
- When implementing account lockout and account policies on Microsoft systems:
    - The Local Security Policy controls policies for user accounts that are defined on a local system.
    - Policy settings in Group Policy are linked to the domain control settings for all user accounts in the domain. Settings defined at other levels in Group Policy do not affect password or account lockout settings.
- Disable and/or remove unnecessary accounts installed on the operating system by default, or disable specific user accounts that are no longer needed.
- Prohibit the use of generic user accounts. Generic accounts, such as guest or administrator accounts in Windows, should be disabled.
- Prohibit the use of shared user accounts.
  Shared accounts:
    - Increase the likelihood of the account being compromised. Because the account is shared, users tend to take security for the account less seriously. For example, one organization found that the passwords for shared user accounts proliferated to the point where hundreds of current and former employees knew them.
    - Make password management more difficult. Because password changes must be communicated to multiple users, many system administrators avoid making any password changes at all. If the password is well known, employees (including former employees that no longer need access to the account) may still know the password.
    - Reduce responsibility for the account. Because users view the account as communal, users may do things with the account that they would not do with their personal account.
    - Destroy audit trails for the account. Because multiple users are associated with the account, it can be difficult to identify who is actually responsible for actions performed with the account.

- Make it difficult to monitor the account for unusual activity. Because multiple users are associated with the account, it is much more difficult to define behaviors that are normal and behaviors that are abnormal. This is problematic because identifying abnormal account activity is key to detecting attacks on your systems.

4.5.11 Configure Smart Card Authentication (Demo Video)

Transcript:

In this demonstration, we'll look at using Group Policy to enforce smart card authentication on computers in a domain. We'll do this by making it so all computers in the research and development department and all administrators need to have a smart card in order to log on.

So, first, let's start by working on the first group, the research and development employees. So, in Server Manager, we're going to go to Tools. We're going to open Group Policy Management. We'll go here, to our Research and Development OU. You can see, in this domain, that we already have that OU created.
Notice that we've already created a Group Policy object and linked it to the Research and Development department OU. So, we'll right-click this Group Policy object and click Edit to open Group Policy Management Editor.
Now, we need to enable two different policies in this Group Policy object to enforce smart card authentication. First, we need to locate a policy named Interactive logon that requires a smart card and enable it. After that, we need to set the interactive log on smart card removal behavior to force log off.
So, let us enable the first policy. Because these are computer policies and not user policies, we'll expand Policies under Computer Configuration. We'll navigate to Windows > Security Settings > Local Policies and then Security Options. Now we need to scroll down to Interactive logon session and locate the policy titled Require Windows Hello for Business or smart card. Keep in mind that because we are editing settings within the Computer Configuration portion of this GPO, the settings will be applied to computers and not to users. If a user in this OU logs onto a computer outside the OU where this policy is applied, they'll not need a smart card.
We're going to double-click the policy. First, we'll define this policy and then enable it. If we click OK, we have just specified that any machines within Research and Development will require a smart card to log in.
The next policy we need to look at is called Interactive logon: Smart card removal behavior. As you can see, it is the one right below it. Double-click that policy. Now, we define this policy. We select Force Logoff. We'll click OK.
We can close the Group Policy Management Editor at this point. Now, we can look and see if the GPO is enforced or not. If you look right here, you can see that the GPO Enforced is labeled as No, which means that the GPO is currently not enforced.
There is no little lock icon, meaning the policy isn't enforced. So, when a GPO is enforced, it means that the settings we specify here in the GPO can't be overruled by another GPO that may be linked to a lower organizational unit within the OU, where the GPO is linked.
So, if we don't enforce it, another GPO, like the one in this sub-OU, could potentially override the settings we defined in our RND GPO right here. Since we don't want that to happen, we want this GPO to be enforced.
We'll right-click on the GPO and select Enforced. Now, the GPO is enforced. Now, no matter what we do at lower layers within this OU, this GPO will always be enforced. The next time someone logs into a computer within Research and Development, they'll have to have a smart card to do so.
If they have their smart card, then they can use it to authenticate. And if they remove the smart card, they'll automatically be forced to log off. So, let's do the same thing for our administrative users. We'll do this through Active Directory Computers and Users. Let's go ahead and close the Group Policy Management window. We'll go to Tools and then Active Directory Users and Computers.
Now we'll go to our Admins OU I've created. Within the Admins OU are all of our administrative users in the domain. We want these users to be forced to use smart card authentication. So, let's click the first user, and we'll highlight both users. If we right-click both of these at the same time, as they are both highlighted, it will open the Properties for all of the users at once.
If you have four or five users, it can be really helpful, so you don't have to go through and do it to each account. Now, we'll click on the Account tab. You can see, down here, that we're going to scroll down to Smart card is required for interactive logon. Select that.

If we click OK, that change is made for all the admin users here. There's one more account we need to do this for, and that is the Default Administrator User account created when Active Directory was initially installed. We'll go into our Users OU.
You can see the default account here. We'll right-click, go to Properties, go to Account, and we can scroll down until Smart card is required for interactive login. We're going to hit OK. Now, when any of these admin users try to authenticate to their local machine using a domain account, they'll be required to have a smart card for authentication.
That's it for this demonstration. In this demo, we talked about enforcing the use of a smart card for authentication to the local computer systems in our domain. We first edited the Group Policy associated with an OU in the domain and required smart card authentication for any computers in that OU. We made it so that if they remove their smart card, they'll be forced to log off immediately. We then required a smart card for user authentication for all administrative users in the domain.

4.5.12 Configure Smart Card Authentication (Simulation)

Scenario

You work as the IT administrator for a growing corporate network. The Research and Development Department is working on product enhancements. Last year, some secret product plans were compromised. As a result, the company decided to implement smart cards for logging on to every computer in the Research and Development Department. No user should be able to log onto the workstation without using a smart card.

In this lab, your task is to perform the following on CorpDC:

- Enforce the existing Research-DevGPO linked to the Research-Dev OU.
- Edit the Research-DevGPO and configure the following local security setting policies located in the Computer Configuration section:

Policy	Setting
Interactive logon: Require smart card	**Enabled**
Interactive logon: Smart card removal behavior	**Force logoff**

Certificate auto-enrollment has already been enabled for the domain.

Explanation

While completing this lab, use the following information to configure the following Security Options policies:

Policy	Setting
Interactive logon: Require smart card	**Enabled**
Interactive logon: Smart card removal behavior	**Force logoff**

Complete this lab as follows:

1. Access the **CorpDC** server.
    a. In Hyper-V Manager, select **CORPSERVER** .
    b. Double-click **CorpDC** .
2. Enforce the existing Research-DevGPO.
    a. From Server Manager, select **Tools** > **Group Policy Management** .
    b. Maximize the window for better viewing.
    c. From the left pane, expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** > **Group Policy Objects** .
    d. From the left pane, select the **Research-DevGPO** .
    e. From the Scope tab, under Links, right-click **Research-Dev** and then select **Enforced** .
3. Edit Research-DevGPO policies.
    a. From the left pane, right-click **Research-DevGPO** and then select **Edit** .
    b. Maximize the window for better viewing.
    c. Under Computer Configuration, expand **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** .
    d. Select **Security Options** .
    e. From the right pane, double-click the ***policy*** .
    f. Select **Define this policy setting** .
    g. Select additional ***parameters*** to configure the policy setting.
    h. Select **OK** .
    i. Repeat steps 3e-3h to configure the additional policy setting.

4.5.13 Smart Card Authentication Facts

This lesson covers the following topics:

- Smart cards
- Smart card benefits and weaknesses

Smart Cards

Smart cards are plastic cards similar to credit cards that have an embedded memory chip that contains encrypted authentication information. Be aware that smart cards:

- Use public key infrastructure (PKI) technology to store digital signatures, cryptography keys, and identification codes.
- Can authenticate a user when used in conjunction with a smart card reader connected to a computer system.
- Typically have RAM, ROM, programmable ROM, and a microprocessor integrated within the card itself.
- Have their own processor, allowing the card to perform its own cryptographic functions.
- Use a serial interface to connect to the card reader.
- Are powered externally by the smart card reader.
- Are generally considered to be tamper-proof.
- Can be divided into two categories:
    - Contact smart cards: these cards use a gold-plated contact pad that must physically touch the contact pad on a smart card reader.
    - Contactless smart cards: these cards do not require physical contact with the reader device. Instead, these cards use Radio Frequency Identification (RFID) technology to communicate with the smart card reader. An antenna is wound around the edge of the card and activated when the card is within proximity of the card reader.

Smart Card Benefits and Weaknesses

Key benefits of smart cards include the following:

- They provide tamper-resistant storage for a user's private key and other personally identifying information (PII).
- They isolate security-related operations from the rest of the system.
- They allow security credentials to be portable.

Smart cards are subject to the following weaknesses:

- Microprobing - this is the process of accessing the chip surface directly to observe, manipulate, and interfere with the circuit.
- Software attacks - these exploit vulnerabilities in the card's protocols or encryption methods.
- Eavesdropping - this captures transmission data produced by the card as it is used.
- Fault generation - this deliberately induces malfunctions in the card.

4.5.14 Practice Questions (Section Quiz)

q_harden_auth_complex_01_secp8

You want to make sure that all users have passwords over eight characters in length and that passwords must be changed every 30 days.

What should you do?

Answers:

- ***Configure account policies in Group Policy.**
- Configure account lockout policies in Group Policy.
- Configure expiration settings in user accounts.
- Configure day/time settings in user accounts.

Explanation:

Configure account (password) policies in Group Policy to enforce rules about the composition of passwords, such as minimum length, complexity, and history requirements.

Use account expiration in a user account to disable an account after a specific day.

Use day/time restrictions to prevent login during certain days or hours.

Account lockout disables a user account after a specified number of incorrect login attempts.

q_harden_auth_complex_02_secp8

You are configuring the Local Security Policy of a Windows system. You want to require users to create passwords that are at least ten characters in length. You also want to prevent login after three unsuccessful login attempts.

Which policies should you configure? (Select two.)

Answers:

- ***Minimum password length**
- Enforce password history
- Account lockout duration
- ***Account lockout threshold**
- Maximum password age
- Password must meet complexity requirements

Explanation:

Set the *Minimum password length* policy to require a password equal to or longer than the specified length. Set the *Account lockout threshold* policy to lock an account after the specified number of incorrect login attempts.

The following lists explains the incorrect policy choices for this scenario:

- *Enforce password history* requires users to input a unique (previously unused) password when changing their password. This prevents users from reusing previous passwords.
- *Maximum password age* forces users to change the password after the specified time interval.
- *Password must meet complexity requirements* prevents using passwords that are easy to guess or crack. It forces passwords to include letters, symbols, and numbers, and also requires passwords of at least seven characters. However, you cannot configure a longer password length requirement with this policy.
- *Account lockout duration* determines the length of time the account is disabled (in minutes). When the time period expires, the account is unlocked automatically.

q_harden_auth_complex_03_secp8

You are teaching new users about security and passwords.

Which of the following is the BEST example of a secure password?

Answers:

- 8181952
- Stiles_2031
- JoHnSmITh
- ***T1a73gZ9!**

Explanation:

The most secure password is T1a73gZ9! because it is eight or more characters in length and combines uppercase and lowercase characters, special symbols, and numbers.

The least secure password is 8181952 because it appears to be a birthday. JoHnSmITh is not secure because it is still a name. Stiles_2031 is more secure but not as secure as random numbers and letters.

q_harden_auth_disable_accounts_secp8

You are the IT Security Manager for a multinational corporation. The company is undergoing a major restructuring, which includes employee layoffs, role changes, and new hires.

Given the situation, which of the following account maintenance practices would be the MOST effective in ensuring the security of your systems?

Answers:

- Implementing a policy that requires all employees to change their passwords immediately.
- ***Disabling the accounts of all employees who have been laid off and reviewing the access rights of remaining employees.**
- Requiring all employees to undergo a new round of security training.
- Implementing a policy that requires all employees to use a password manager.

Explanation:

Disabling the accounts of all employees who have been laid off is crucial to prevent unauthorized access and the most effective in this scenario. Reviewing the access rights of remaining employees is also important, as their roles may have changed during the restructuring. This option directly addresses the issue at hand and is the most effective way to maintain account security in this situation.

While requiring all employees to change their passwords immediately might seem like a good idea, it could lead to confusion and potentially weak password choices. Moreover, it doesn't address the issue of access rights for employees who have been laid off or whose roles have changed.

While security training is important, it's not the most effective immediate response to the situation. Training takes time and doesn't address the immediate risk of unauthorized access from laid-off employees or employees with changed roles.

While using a password manager is a good security practice, it doesn't address the immediate risks associated with the company's restructuring. It won't prevent unauthorized access from accounts that belong to laid-off employees or employees whose roles and access rights have changed.

q_harden_auth_history_secp8

You are configuring the Local Security Policy of a Windows system. You want to prevent users from reusing old passwords. You also want to force them to use a new password for at least five days before changing it again.

Which policies should you configure? (Select two.)

Answers:

- ***Enforce password history**
- Maximum password age
- ***Minimum password age**
- Password must meet complexity requirements

Explanation:

Set the *Enforce password history* policy to prevent users from reusing old passwords. Set the *Minimum password age* policy to prevent users from changing passwords too soon. Passwords must remain the same for at least the time period specified.

Use the *Maximum password age* policy to force periodic changes to the password. After the maximum password age has been reached, the user must change the password.

Use the *Password must meet complexity requirements* policy to require that passwords include letters, numbers, and symbols. This makes it harder for hackers to guess or crack passwords.

q_harden_auth_lockout_secp8

For users on your network, you want to automatically lock user accounts if four incorrect passwords are used within ten minutes.

What should you do?

Answers:

- Configure account expiration in user accounts.
- Configure day/time restrictions in user accounts.
- ***Configure account lockout policies in Group Policy.**
- Configure password policies in Group Policy.
- Configure the enable/disable feature in user accounts.

Explanation:

Account lockout disables a user account after a specified number of incorrect login attempts. The account lockout threshold identifies the allowed number of incorrect login attempts. The account lockout counter identifies a time period for keeping track of incorrect attempts (such as 10 minutes).

If account lockout locks a user account, use the unlock feature to allow login. Use the enable/disable feature to prevent or allow login using the user account.

Configure account (password) policies in Group Policy to enforce rules about the composition of passwords, such as minimum length, complexity, and history requirements. Use account expiration in a user account to disable an account after a specific day. Use day/time restrictions to prevent login during certain days or hours.

q_harden_auth_multifactor_secp8

You are the IT security manager for a rapidly growing tech company. The company has been using simple password authentication for all systems.

However, with the increasing number of employees and the sensitivity of the data being handled, you decide it's time to harden the authentication methods.

Which of the following steps would be the MOST effective in achieving this goal?

Answers:

- Implementing a policy that requires all passwords to be at least 8 characters long.
- ***Implementing multifactor authentication (MFA) for all systems.**
- Requiring all employees to change their passwords every 30 days.
- Implementing a policy that allows employees to use their personal email addresses for system logins.

Explanation:

Implementing multifactor authentication (MFA) is the most effective option. MFA requires users to provide two or more verification factors to gain access to a resource such as an application, online account, or a VPN. It adds an additional layer of security, reducing the likelihood of successful cyber attacks.

While having a policy that requires all passwords to be at least 8 characters long is a step in the right direction, it is not the most effective way to harden authentication methods. Longer passwords are generally more secure, but they can still be compromised, especially if they are not complex.

Requiring all employees to change their passwords every 30 days can lead to password fatigue, where users resort to using weak or reused passwords. While regular password changes can be part of a secure system, it is not as effective as implementing MFA.

Allowing employees to use their personal email addresses for system logins is not a good practice for hardening authentication. Personal email addresses can be compromised, which could lead to unauthorized access to company systems. It's better to use work-specific email addresses or other forms of identification for system logins.

q_harden_auth_password_expiration_secp8

The IT department at a small company is revamping its password policies to bolster security. The company wants to ensure employees follow best practices for creating and managing passwords.

The department aims to promote a secure environment by implementing password expiration policies.

Which method for password management is BEST to promote a secure environment by requiring users to change their passwords after a certain period?

Answers:

- ***Password expiration**
- Password complexity
- Password reuse prevention
- Password recovery via email

Explanation:

Implementing a password expiration policy requires users to change their passwords after a set period. This practice helps reduce the risk of unauthorized access from compromised passwords obtained in the past.

Password complexity requirements are essential for creating strong passwords, but they do not directly address the impact of password expiration on overall security.

Preventing password reuse is crucial to avoid using the same password for multiple accounts, but it is not directly related to the role of password expiration in enhancing security.

Password recovery via email involves users recovering forgotten passwords through email verification.

q_harden_auth_password_history_secp8

An employee at a company frequently recycles old passwords when prompted for a password change.

What feature of a password policy can prevent this?

Answers:

- Password length
- Password complexity
- Password age
- ***Password history**

Explanation:

The password history attribute keeps track of previously used passwords and prevents employee from using them again, discouraging password recycling.

Password length refers to the minimum (and sometimes maximum) length that a password should have. It does not prevent the reuse of old passwords.

Password complexity enforces rules like a mix of characters, numbers, and symbols. It does not prevent the reuse of old passwords.

Password age is the length of time an employee can use a password before it expires, and the employee must choose a new one. It does not prevent the reuse of old passwords.

q_harden_auth_reuse_secp8

You have just configured the password policy and set the minimum password age to 10.

What is the effect of this configuration?

Answers:

- ***Users cannot change the password for 10 days.**
- Users must change the password at least every 10 days.
- The password must contain 10 or more characters.
- The password must be entered within 10 minutes of the login prompt being displayed.
- The previous 10 passwords cannot be reused.

Explanation:

The minimum password age setting prevents users from changing the password too frequently. After the password is changed, it cannot be changed again for at least 10 days.

The maximum password age setting determines how frequently a password must be changed.

The minimum password length setting controls the minimum number of characters that must be in the password.

Password history is used to prevent previous passwords from being reused.

q_harden_auth_rshared_secp8

Upon running a security audit in your organization, you discover that several sales employees are using the same domain user account to log in and update the company's customer database.

Which action should you take? (Select two. Each response is part of a complete solution.)

Answers:

- ***Delete the account that the sales employees are currently using.**
- ***Train sales employees to use their own user accounts to update the customer database.**
- Implement a Group Policy Object (GPO) that restricts simultaneous logins to one.
- Implement a Group Policy Object (GPO) that implements time-of-day login restrictions.
- Apply the Group Policy Object (GPO) to the container where the sales user accounts reside.

Explanation:

You should prohibit the use of shared user accounts. Allowing multiple users to share an account increases the likelihood of the account being compromised. Because the account is shared, users tend to take security for the account less seriously. In the scenario, the following tasks need to be completed:

- The existing shared user account needs to be deleted. Until you delete the account, users can continue to use it for authentication. You could just change the password on the account, but there is a high chance that the new password would be shared again.
- Train sales employees to use their own user accounts to update the customer database. Ensure that these accounts have the level of access required for users to access the database.

Applying time-of-day login restrictions in a Group Policy object does not address the issue in this scenario.

q_harden_auth_shared_accounts_secp8

You are the IT security manager for a large corporation. The company has been using shared accounts for certain systems due to ease of access and convenience.

However, you are considering implementing a policy to prohibit the use of shared accounts.

Which of the following are valid reasons for this decision? (Select two.)

Answers:

- Shared accounts allow for easier password management.
- ***Shared accounts can lead to accountability issues.**
- Shared accounts reduce the need for individual user training.
- ***Shared accounts can compromise the principle of least privilege.**
- Shared accounts increase the speed of system access.

Explanation:

The following are valid reasons for prohibiting the use of shared accounts:

- Shared accounts can lead to accountability issues, making it difficult to track who did what. If an issue arises, it's nearly impossible to hold the appropriate person accountable because multiple people have access to the same account.
- The principle of least privilege states that users should only have access to the resources they need to do their jobs and nothing more. Shared accounts often have broad access rights, which can lead to unauthorized access to sensitive information.

Shared accounts may seem easier to manage because there is only one set of credentials to remember. However, they actually complicate password management because when one person leaves the company or changes roles, the password must be changed and communicated to all other users.

While it might seem that shared accounts reduce the need for individual user training, this is not a valid reason to use them. Each user should be trained on the systems they need to use to perform their job effectively and securely.

Shared accounts may seem to increase the speed of system access, but this is a short-term benefit. In the long term, they can lead to serious security risks, which can cause significant delays and costs if a breach occurs.

q_harden_auth_time_secp8

You have hired ten new temporary employees to be with the company for three months.

How can you make sure that these users can only log on during regular business hours?

Answers:

- ***Configure day/time restrictions in user accounts.**
- Configure account expiration in user accounts.
- Configure account policies in Group Policy.
- Configure account lockout in Group Policy.

Explanation:

Use day/time restrictions to limit the days and hours when users can log on.

Configure account expiration to disable an account after a specific date.

Use account policies in Group Policy to configure requirements for passwords.

Use account lockout settings in Group Policy to automatically lock accounts when a specific number of incorrect passwords are entered.

q_smartcard_auth_attack_secp8

Match each smart card attack on the left with the appropriate description on the right.

Answers:

- Accesses the chip's surface directly to observe, manipulate, and interfere with a circuit.
- Exploits vulnerabilities in a card's protocols or encryption methods.
- Captures transmission data produced by a card as it is used.
- Deliberately induces malfunctions in a card.

Explanation:

Smart cards are subject to the following weaknesses:

- Microprobing is the process of accessing a chip's surface directly to observe, manipulate, and interfere with the circuit.
- Software attacks exploit vulnerabilities in the card's protocols or encryption methods.
- Eavesdropping captures transmission data produced by the card as it is used.
- Fault generation deliberately induces malfunctions in a card.

q_smartcard_auth_contactless_secp8

John, a security analyst, is using a smart card to gain access to a secure server room. He simply waves his card near the card reader and the door unlocks.

Later, he uses the same card to log into his computer by inserting it into a card reader.

Based on this information, is John using a contact or contactless smart card?

Answers:

- Contact smart card, because he inserted the card into his computer's card reader.
- Contactless smart card, because he waved the card near the door's card reader.
- ***Both a contact and contactless smart card, because he used the card both by inserting it into a reader and by waving it near a reader.**
- Neither a contact nor contactless smart card, because smart cards cannot be used both ways.

Explanation:

John is using a smart card that has both contact and contactless capabilities. This type of card is known as a dual-interface smart card. It can be used by inserting it into a card reader (contact) and by waving it near a reader (contactless).

While John did insert the card into his computer's card reader, which is a characteristic of contact smart cards, this does not account for the fact that he also used the card by waving it near the door's card reader, which is a characteristic of contactless smart cards.

While John did wave the card near the door's card reader, which is a characteristic of contactless smart cards, this does not account for the fact that he also used the card by inserting it into his computer's card reader, which is a characteristic of contact smart cards.

While it may seem unusual, there are smart cards that have both contact and contactless capabilities. These are known as dual-interface smart cards. So, it is possible for a smart card to be used in both ways.

q_smartcard_auth_key_benefits_secp8

Which of the following are key benefits of using smart cards? (Select two.)

Answers:

- ***They provide tamper-resistant storage for a user's private key and other personally identifying information (PII).**
- They allow for unlimited data storage.
- They can be used to exploit vulnerabilities in a system's protocols.
- ***They isolate security-related operations from the rest of the system.**
- They can induce malfunctions in the card reader.

Explanation:

The following are key benefits of using smart cards:

- One of the key benefits of smart cards is that they provide tamper-resistant storage for a user's private key and other personally identifying information (PII). This makes them a secure method of storing sensitive information.
- Smart cards isolate security-related operations from the rest of the system. This means that even if a system is compromised, the operations carried out by the smart card remain secure.

While smart cards do have storage capabilities, they do not allow for unlimited data storage. The storage capacity of a smart card is limited by the size of the embedded memory chip.

Smart cards do not exploit vulnerabilities in a system's protocols. In fact, they are designed to enhance security by isolating security-related operations from the rest of the system.

Smart cards do not induce malfunctions in the card reader. They are designed to work seamlessly with card readers and any malfunctions would likely be due to a fault in the card reader itself, not the smart card.

q_smartcard_auth_key_weaknesses_secp8

Which of the following are key weaknesses of using smart cards? (Select two.)

Answers:

- ***They are susceptible to software attacks that exploit vulnerabilities in the card's protocols or encryption methods.**
- They are unable to store digital signatures, cryptography keys, and identification codes.
- ***They are vulnerable to eavesdropping that captures transmission data produced by the card as it is used.**
- They are incapable of performing their own cryptographic functions.
- They require a constant power supply to function.

Explanation:

The following are key weaknesses of using smart cards:

- One of the key weaknesses of smart cards is that they are susceptible to software attacks that exploit vulnerabilities in the card's protocols or encryption methods. This means that if a hacker can find a weakness in the card's security protocols, they may be able to gain unauthorized access to the card's data.
- Another weakness of smart cards is that they are vulnerable to eavesdropping that captures transmission data produced by the card as it is used. This means that if a hacker can intercept the data being transmitted by the card, they may be able to gain unauthorized access to the card's data.

Smart cards are actually capable of storing digital signatures, cryptography keys, and identification codes. This is one of their key features and not a weakness.

Smart cards are actually capable of performing their own cryptographic functions. This is one of their key features and not a weakness.

While smart cards do require power to function, they do not require a constant power supply. They are powered externally by the smart card reader when they are inserted into the reader. This is not considered a weakness of smart cards.

q_smartcard_auth_pki_secp8

Which technology is primarily used by smart cards to store digital signatures, cryptography keys, and identification codes?

Answers:

- Blockchain technology
- ***Public Key Infrastructure (PKI)**
- Secure Sockets Layer (SSL)
- Advanced Encryption Standard (AES)
- Hashing algorithms

Explanation:

Public Key Infrastructure (PKI) is the technology primarily used by smart cards. It allows for the storage of digital signatures, cryptography keys, and identification codes, providing secure and encrypted communication.

Blockchain technology is primarily used for creating decentralized digital ledgers for transactions. It is not typically used in the context of smart cards for storing digital signatures, cryptography keys, and identification codes.

Secure Sockets Layer (SSL) is a protocol used for securing data transmission over the internet. While it does involve encryption and can work with PKI, it is not the primary technology used by smart cards for storing digital signatures and keys.

Advanced Encryption Standard (AES) is an encryption standard used to secure data. While it may be used in conjunction with PKI for encryption purposes, it is not the primary technology used by smart cards for storing digital signatures and keys.

Hashing algorithms are used to convert data into a fixed size of numerical or alphanumeric characters. While they play a role in encryption and security, they are not the primary technology used by smart cards for storing digital signatures and keys.

4.6 Linux Users

As you study this section, answer the following questions:

- How do you create a user in Linux?
- Why shouldn't passwords expire too frequently?
- Which directory contains configuration file templates copied into a new user's home directory?
- Which command deletes a user and the user's home directory simultaneously?

In this section, you will learn to:

- Create a user account
- Rename a user account
- Delete a user
- Change your password
- Change a user's password
- Lock and unlock user accounts
- Configure Linux User Security and Restrictions
- Configure SELinux

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise. • Hardening techniques ○ Default password changes 4.1 Given a scenario, apply common security techniques to computing resources. • Application security 4.5 Given a scenario, modify enterprise capabilities to enhance security.

Exam	Objective

	• Operating system security
	• SELinux
	
	4.6 Given a scenario, implement and maintain identity and access management.
	
	• Provisioning/de-provisioning user accounts
	• Password concepts
	  ○ Reuse
	  ○ Expiration

TestOut Security Pro	1.1 Manage identity 1.1.2 Manage Linux users and groups 1.2 Harden authentication 1.2.2 Manage account password

4.6.1 Linux User and Group Overview (Lesson Video)

Transcript:

In this video, we're going to discuss Linux users and groups. First, when you're working with Linux, it's important to understand that it's a true multi-user operating system. Each user account has their own unique, customizable computing environment. When user accounts are stored locally, Linux stores the user, group, and password information in three separate files on the local file system in the /etc directory. We'll discuss each of these files in this lesson.

The first file is the passwd file. The password file is named passwd, and it's stored in the /etc directory. This file contains the user account information for your system; this is where your users are defined. But don't let the filename confuse you. This file doesn't contain any password information, only user account information. Passwords are stored in another file that we'll look at in a bit.
You may hear people still refer to this file as the password file because in the early days of UNIX, passwords were stored in here alongside the usernames. But, as you can imagine, this presented a huge security risk, so this file no longer contains any passwords. To avoid confusion, I'll call this the passwd file.
So, now, the passwd file is where all the user accounts are stored. Each user account on a system is represented by a single line called an account record. An account record is composed of several fields, each separated by a colon.
The first field is the username. The second is the password field. The third is the user ID field. Then there's the group ID field, the full name field, the home directory field, and finally, the shell field.
Let's look at an example of each field. The username field contains the username the user will use to log in to the system. For example, for this record, the user account is zday. Next, we have the password field. Since this is a remnant of when passwords used to be stored in this file, you'll only see an X in this field. Now passwords are stored in the shadow file, which we'll talk about in a bit. Next is the UID field. This contains the user ID that's assigned to the user account. This user account is assigned a UID number of 1001.
Then we have the group ID field, which references the group ID number of the user's default group. In Linux, a single user account can be a member of several different groups, but by default, a user is always assigned to one group, the default group. This is almost always the Users group, which has an ID of 100.
Next, we have the full name field, which just contains the user's full name. The next field specifies the user's home directory. And finally, the default shell field specifies which shell is going to run for the user by default. Most distributions set the default shell to /bin/bash.

It's important to understand that the passwd file contains two different types of users. First, there's the standard user accounts. These accounts are used to log in to the system, such as the zday account we just saw.
The second type of accounts are the system user accounts. These accounts can't be used to log in to the system. They're used by services or daemons that are running on the system. This is because Linux uses groups, users, and permissions to control privileges. For a service or daemon to be able to read from or
write to the file system, it needs to have a user account. So, if you ever open the passwd file and wonder why there are so many users that you never created, it's because the passwd file contains both standard user accounts.
Now let's talk about the shadow file, which is where passwords are stored these days.
The shadow file is linked to the passwd file. Each user account defined in your passwd file has a corresponding entry in the shadow file. As with the passwd file, each user account is represented by a single line in the shadow file called a record, which is composed of several different fields separated by colons.
Here's an example of a single record in the shadow file.
First, we have the username field, which is the same username that's defined in the passwd file. This is what links the two records together. Next, we have the password field. Depending on the type of user account you're looking at, there are a few things you can see in the password field.
For example, if the record is for a standard user account, the password field will contain a string of text. This is the user's actual password in encrypted format. Theoretically, it's possible to crack this encrypted password if someone were to gain access to this file. But it's extremely difficult to do, and it would take a very, very long time.
If the record is for a system user account, the password field will contain an asterisk. Remember, system user accounts can't be used to log in to the system. The asterisk indicates that this account is disabled, which prevents it from being able to authenticate to the system.
There's one more thing you might see in the password field: a blank entry. Nothing. This means that the user account doesn't have a password and doesn't require one to log in to the system. As you can imagine, this presents a huge security hole. A user account should always have either an encrypted password or an asterisk, not a blank entry.
The next field is the last modified field. This field displays the numbers of days since January 1st, 1970 that the password was last changed. So, when you see a huge number here, just remember that it's counting all the days between January 1st, 1970, and the day the password was last changed.
Next, we have the minimum days field. This field defines the minimum number of days that need to pass before a password can change. Usually, this is set to zero, which means the password can be changed any time. But if you want to make it so users need to keep the same password for at least a week, you could put a seven in this field.
The next field is the maximum days field. This field defines the maximum number of days a password can be used. For example, if you want users to change their password after 90 days, you would enter '90' in this field, or you can set passwords to never expire by entering â€˜99999' here.
Next, we have the days warning field. This field defines when a user will be warned that their password is about to expire. This field is usually set to five or seven days.
The next field defines the number of days to wait after a password is expired to disable the account. This field is usually set to -1, which disables this functionality. But if you want to disable an account if its password isn't changed after it expires, then you can enter a number, 3 for example.
And finally, we have the expire field. This field defines the number of days since January 1st, 1970, that must pass until the account is automatically disabled. By default, this field is blank, which means this functionality is disabled. But if you want an account to be automatically disabled on a specific date, you need to figure out the number of days between it and January 1st, 1970, and then enter that number. And don't worry, there are several tools available that can do this calculation for you.
Now, because the passwd and shadow files work together, they need to stay synchronized. The main way these files become out of sync is when an administrator uses a text editor to make changes to one of the files, but forgets to make the changes to the other, such as changing a username. If the files don't match, you've got a problem.
To avoid this, manage user accounts using the password and user management utilities that are provided by the operating system. Avoid using a text editor. If a problem does occur, there are few commands you can use to compare the passwd file and shadow file, such as pwck, and then re-synchronize the files if they don't match.
Now, let's look at the third file, the /etc/group file. Like most operating systems,
Linux uses groups to help ease system management. For example, instead of managing the privileges of each individual user, you can manage the privileges of a single group and then assign user accounts to that group.
Now, any changes made to the group will be automatically applied to the users. If you're managing hundreds of users, groups make your job a lot easier. Groups on a Linux system are defined in the /etc/group file. And the structure of the group file is similar to the passwd and shadow files.
Each line in the group file is a single record that defines one group on your system, and each record is composed of four fields separated by colons. The first field is the group field. It specifies the group's name. In this example, the name of

the group is sales. The second field is the password field, which specifies a password for the group. This field is rarely used, and in most cases it'll contain an X, which means no password is set. The third field is the group ID number, or the GID. In this case, the GID is 33. And the last field is a list of users who are members of the group. In this example, you see that zday and rnelson are both members of the sales group.
Keep in mind that some Linux distributions use an additional group file to store group passwords. This works in a similar way to the passwd and shadow files. With this model, group passwords are stored in the gshadow file. Like the group file, the gshadow file has corresponding records and fields that define each group. In the gshadow file, you'll see fields that define the group name, the password, the group admins, and the group members.
So that's how Linux users and groups work on a Linux system. In this lesson, we talked about user account storage. First, we looked at the /etc/passwd file, which contains user account information. Then we looked at the shadow file, which contains user passwords and password settings. And finally, we looked at the group file, where groups and group members are defined.

4.6.2 Managing Linux Users (Demo Video)

Transcript:

In this demonstration, we're going toÂ talk about managing Linux user accounts.Â As a Linux administrator, there will be times whenÂ you have to add, remove, or modify user accounts.Â Let's begin by looking at the User AddÂ utility.

As a root user, you type â€˜useradd' followed byÂ any options you might want and then the username.Â It's important to note that if you create a userÂ in the shell prompt with useradd, it will automatically use the default valuesÂ contained and the /etc/default/user add file.
Let's look at that file right now withÂ 'less /etc/default/useradd'.Â Notice that the useradd commands use the default values. First, look at the defaultÂ home directory, specified as the /home directory.
That's where the user's default profile, or default directory,Â will be created. There's no expiration date set by default for the account.Â The shell is bin/bash, and theÂ skeleton file is etc/skel.
This is useful if you have specific files or foldersÂ that you want to be in each new user's home directory.Â You can simply place them in the skeletonÂ directory, and then they're automaticallyÂ created in the new user's home directory.
We can see here, at the bottom, that a mailÂ spool for the user account will also be created.Â You can go into this file usingÂ the VI editor and change these default values if you need to.
For example, if you want the home directory to be inÂ a different location by default, you can set thatÂ default by typing â€˜useradd -D' at the shell prompt.Â You need to also be aware that the less /etc/login.defsÂ file is used to configure the values that can beÂ used for the group ID and the user ID.
If we hit Enter and scroll down,Â we can see here that the values for the group ID andÂ the user ID are specified, so you can see the minimumÂ value for the user ID on this Linux version is 1000.Â Also, for the group ID, the minimum is set to 1000.Â Let's exit out of that.
Now let's talk about useradd.Â Like I said before, you can override theÂ default values using command line options. We'll take a look at the man page for useradd, â€˜man useradd'. This will show you all the available options for the useradd command. A few examples are -c to specify a comment for the user's full name, -g (lowercase) to add a single primary group, and -G (uppercase) to add supplementary groups separated by commas.
Let's look at an example.Â Suppose we want to create a user namedÂ ksanders whose full name is Kim Sanders.Â We want to create a home directory, and we want toÂ specify the user's name as we create the account.Â Let's go ahead and type in 'sudo useradd'.
Now we want to use the -c optionÂ to specify the user's full name.Â We're going to put in â€˜"Kim Sanders"' in quotes.Â Then we're going to use the â€˜-m' option to create the user'sÂ home directory and specify the username as â€˜ksanders'.Â Press Enter.
After suppling the sudo password to add the account, we can verify that was created by looking at our etc/passwd file.Â If we scroll down, we can see a few things here.Â You can see the group and user ID, 1002, 1002.We can see the full name, Kim Sanders, and we can alsoÂ see the ksanders folder within the home directory.
ksanders doesn't have a password yet.Â On Linux, if the user account doesn't haveÂ a password, then the password is locked.Â We can verify this by typing in 'sudo passwdÂ -S' followed by the username, 'ksanders'.
Now we can see this LK, which tells usÂ that the account is locked.Â So, before the user can log in, we have toÂ enable the account by adding a password.

We need to type 'sudo passwd'Â followed by the user account, 'ksanders'.Â Hit Enter.Â It wants us to create a password.Â We're going to type '1234' just for now. This is just a test account, so that password is fine.Â Retype the password, '1234'.Â We get a warning about the password length, but the password updates successfully.Â Now we want to verify that, so we'reÂ going to retype â€˜sudo passwd-S ksanders'.
You can see now that this has changed to a PS from an LK.Â That indicates that a passwordÂ has been unlocked and is now set up for this account.Â Now this user can actually log in to this profile.
From time to time, you may need to modify an existing user.Â This is done with the usermod command. usermod's options are prettyÂ much the same as useradd's.
For example, let's say we entered theÂ wrong full name for the ksanders account.Â Her name is actually Kimberly.Â We can go ahead and change this by typing 'sudoÂ usermod - c' and then putting the correct name, 'Kimberly Sanders',Â closed single quotations, and then the user account, 'ksanders'.
Now we want to verify those changes.Â We're going to go to 'lessÂ /etc/passwd' and scroll to the bottom.Â We can see that 'Kim Sanders'Â has actually changed to 'Kimberly Sanders'.
You can use other options with the usermod command to changeÂ user account information, such as the username itself, theÂ user ID, the default group, and the home directory.
The last thing we want to look at isÂ the user delete, or userdel, command.Â This command is used to remove an account.Â We're going to remove the ksanders account.
There's only one option we can use withÂ userdel, and that's the -r option.Â By default, if you don't use the -r option,Â the userdel command won't delete theÂ user's home directory from the file system. This is because there might be intellectualÂ property in that directory that you need toÂ keep, even if you delete the user account.Â But if you do want to delete the home directory alongÂ with the user account, you can use the -r option. Let's go ahead and type this in, 'sudo userdelÂ ksanders -r'. We're not going to hit Enter justÂ yet because I want to show you another way of addingÂ and maintaining user accounts.
Sometimes, depending on the distribution, you can actually do this through aÂ GUI or graphical user interface within Linux.Â Some Linux systems have them, and some don't.Â Sometimes you can download them.Â This system already has one. Let's go ahead and click on the dropdown by the power button and click Settings. Scroll down to Details and click Users.Â Here's our Kim Sanders account. It looks like the name didn't update to Kimberly in the GUI, but we do know it changed because we just saw that in the command line.Â If I were to unlock this, I could go in andÂ manage this account.
Here, I can do things such change the account type or reset the password.Â It'll show me my login history as well.Â I'm not actually going to do anything from here,Â but I did want to show you that because I want toÂ show you how the account is actually removed.
If I go ahead and I run thisÂ command, now the user is deleted.Â If we go back and look at the â€˜/etc/passwd' fileÂ and scroll down, you can see the account isÂ deleted.
That's it for this demonstration.Â In this demo, we talked about managingÂ Linux users using the shell prompt.Â We looked at the useradd command, the usermodÂ command, the passwd command, and the userdel command.
We also looked at the graphic interface thatÂ some Linux systems will use, and we looked atÂ other options for each command.

4.6.3 Linux User Commands and Files

This lesson covers the following topics:

- User files
- User management commands

User Files

Linux is highly flexible regarding where user and group information is stored. The options for storing the data are:

- Local file system.
- LDAP-compliant database.
- Network Information System (NIS). NIS allows many Linux computers to share common user accounts, group accounts, and passwords.

- A Windows domain.

When files are stored in the local file system, the following files are used:

File	Description
/etc/passwd	The /etc/passwd file contains the user account information. Each user's data is stored in a single line on this file. There are two types of accounts in a Linux system: • Standard accounts (these are user accounts). • System user accounts (these are used by services).
/etc/shadow	In Linux, local user account names are stored in /etc/passwd. When a user logs in to a local interactive shell, the password is checked against a hash stored in /etc/shadow. There are corresponding entries in both files, and they must stay synchronized. The system provides password and user management utilities, allowing you to edit and keep the files synchronized. You can use the following commands to identify errors and synchronize the files: • **pwck** verifies each line in the two files and identifies discrepancies. • **pwconv** adds the necessary information to synchronize the files. Interactive login over a network is typically accomplished using Secure Shell (SSH). With SSH, the user can be authenticated using cryptographic keys instead of a password. A pluggable authentication module (PAM) is a package for enabling different authentication providers, such as smart-card log-in. The PAM framework can also be used to implement authentication to network directory services.
/etc/group	As with Active Directory, groups can be used to simplify user access to network resources. The /etc/group file contains information about each group.

Be aware of the following configuration files when managing user accounts:

File	Description
/etc/default/useradd	The /etc/default/useradd file contains default values used by the useradd utility when creating a user account, including: • Group ID. • Home directory. • Account expiration. • Default shell. • Secondary group membership.
/etc/login.defs	The /etc/login.defs file contains: • Values used for the group and user ID numbers. • Parameters for password encryption in the shadow file. • Password expiration values for user accounts.

File	Description
/etc/skel	The /etc/skel directory contains a set of configuration file templates that are copied into a new user's home directory when it is created, including the following files: • .bashrc • .bash_logout • .bash_profile • .kshrc

User Management Commands

Although it is possible to edit the /etc/passwd and /etc/shadow files manually to manage user accounts, doing so can disable your system. Instead, use the following commands to manage user accounts:

Command	Command Function
useradd	Create a user account. The following options override the settings as found in /etc/default/useradd: • **-c** adds a description for the account in the GECOS field of /etc/passwd. • **-d** assigns an absolute pathname to a custom home directory location. • **-D** displays the default values specified in the /etc/default/useradd file. • **-e** specifies the date on which the user account will be disabled. • **-f** specifies the number of days after a password expires until the account is permanently disabled. • **-M** defines the secondary group membership. • **-m** creates the user's home directory (if it does not exist). • **-n** does not create a group with the same name as the user (Red Hat and Fedora, respectively). • **-p** defines the encrypted password. • **-r** specifies that the user account is a system user. • **-s** defines the default shell. • **-u** assigns the user a custom UID. This is useful when assigning ownership of files and directories to a different user.
passwd	Assign or change a password for a user: • **passwd** (without a username or options) changes the current user's password. • Users can change their own passwords. The root user can execute all other **passwd** commands. Be aware of the following options: • **-S** username displays the status of the user account. LK indicates that the user account is locked, and PS indicates the user account has a password. • **-l** disables (locks) an account. This command inserts a !! before the password in the /etc/shadow file, effectively disabling the account. • **-u** enables (unlocks) an account. • **-d** removes the password from an account.

Command	Command Function
	• **-n** sets the minimum days before a password can be changed. • **-x** sets the number of days before a user must change the password (password expiration time). • **-w** sets the number of days before the password expires that the user is warned. • **-t** sets the number of days following the password expiration that the account will be disabled.
usermod	Used to modify an existing user account; usermod uses several of the same switches as useradd. Be aware of the following switches: • **-c** changes the description for the account. • **-l** renames a user account. • **-L** locks the user account. This command inserts a ! before the password in the /etc/shadow file, effectively disabling the account. • **-U** unlocks the user account.
userdel	Remove the user from the system. Be aware of the following options: • **userdel** [username] (without options) removes the user account. • **-r** removes the user's home directory. • **-f** forces removing the user account even when the user is logged into the system.

If you are logged in as the root user, the commands in the table can be run by typing the applicable command and its options. However, if you are not logged in as the root user, you must use the **sudo** or **su** command to gain the permissions required. For example, to create a new user named Kim Sanders, you would run: **sudo useradd -c "Kim Sanders" -m ksanders**

4.6.4 Create a User Account (Simulation)

Scenario

The VP of marketing has told you that Paul Denunzio will join the company as a market analyst in two weeks. You need to create a new user account for him.

You are logged in as root, so the **sudo** command is unnecessary.

In this lab, your task is to:

- Create the **pdenunzio** user account.
  - Include the full name, **Paul Denunzio** , as a comment for the user account.
- Set **eye8cereal** as the password for the user account.
- When you are finished, view the **/etc/passwd** file to verify the creation of the account.
- Answer the question.

Explanation

Complete this lab as follows:

1. Create the Paul Denunzio account and comment.
    a. From the Favorites bar, select **Terminal** .
    b. From the Linux prompt, type **useradd -c "Paul Denunzio" pdenunzio** and press **Enter** .
2. Create a password for Paul.
    a. Type **passwd pdenunzio** and press **Enter** .
    b. Type **eye8cereal** as the password and press **Enter** .
    c. Retype **eye8cereal** as the password and press **Enter** .
3. Verify that the account was created.
    - Type **cat /etc/passwd** and press **Enter** .
4. Answer the question.
    a. From the top right, select **Answer Questions** .
    b. Select the correct answer.
    c. Select **Score Lab** .

4.6.5 Rename a User Account (Simulation)

Scenario

Brenda Cassini (bcassini) was recently married. You need to update her Linux user account to reflect her new last name, Palmer. You are currently logged in as the root account, so you will not need to use the sudo command to get permissions to perform the required tasks.

In this lab, your task is to open Terminal and then use the **usermod** command to:

- Rename Brenda's user account to **bpalmer** .
- Change Brenda's comment field to read **Brenda Palmer** .
- Change and move Brenda's home directory to **/home/bpalmer** .
- When you're finished, view the **/etc/passwd** file and **/home** directory to verify the modification.

Explanation

Complete this lab as follows:

1. Rename the bpalmer account and move her home directory.
    a. From the Favorites bar, select **Terminal** .
    b. From the Terminal prompt, type **usermod -l bpalmer bcassini -m -c "Brenda Palmer" -d /home/bpalmer** and press **Enter** .
2. Verify account modification.
    a. Type **cat /etc/passwd** and press **Enter** .
    b. Find the line that shows that Brenda's account has been changed.
    c. Type **ls /home** and press **Enter** to verify that the account was modified.
    Notice that the home directory for Brenda is now bpalmer.

4.6.6 Delete a User (Simulation)

Scenario

Terry Haslam (thaslam) was dismissed from the organization. His colleagues have harvested the files they need from his home and other directories. Your company security policy states that upon dismissal, the user's accounts should be removed in their entirety.

In this lab, your task is to:

- Delete the thaslam user account and home directory from the system.
- When you're finished, view the **/etc/passwd** file and **/home** directory to verify the account's removal.

Explanation

Complete this lab as follows:

1. Delete the Terry Haslam account and home directory.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **userdel -r thaslam** and press **Enter** .
2. Verify the account's removal.
    a. Type **cat /etc/passwd** and press **Enter** .
    b. Type **ls /home** and press **Enter** to verify the account was removed.

4.6.7 Change Your Password (Simulation)

Scenario

You use a special user account called Administrator to log on to your computer. However, you think someone has learned your password. You are logged on as Administrator.

In this lab, your task is to change your password to **r8ting4str** . The current Administrator account uses **7hevn9jan** as the password.

As you type in the password, the cursor will not move. Continue entering the password anyway.

Explanation

Complete this lab as follows:

1. Change your password.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **passwd** and press **Enter** .
    c. When prompted, enter **7hevn9jan** and press **Enter** . This is the current password.
    d. At the *New password* prompt, enter **r8ting4str** and press **Enter** .
    e. Retype **r8ting4str** as the new password and press **Enter** .

4.6.8 Change a User's Password (Simulation)

Scenario

Salman Chawla (schawla) forgot his password and needs access to the resources on his computer. You are logged on as wadams. The password for the root account is **1worm4b8** .

In this lab, your task is to:

- Change the password for the schawla user account to **G20oly04** (0 is a zero).
- Make sure the password is encrypted in the shadow file.

Do not use the **usermod -p** command to change the password, as this stores the unencrypted version of the password in the /etc/shadow file.

Explanation

Complete this lab as follows:

1. Change Salman Chawla's password.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **su -c "passwd schawla"** , then press **Enter** .
    c. Type **1worm4b8** , then press **Enter** . This is the password for the root user.
    d. At the *New password* prompt, type **G20oly04** , then press **Enter** . This is the new password for the schawla user account.
    e. At the *Retype new password* prompt, type **G20oly04** , then press **Enter** .

4.6.9 Lock and Unlock User Accounts (Simulation)

Scenario

Every seven years, your company provides a six-week sabbatical for every employee. Vera Edwards (vedwards), Corey Flynn (cflynn), and Bhumika Kahn (bkahn) are leaving today. Maggie Brown (mbrown), Brenda Cassini (bcassini), and Arturo Espinoza (aespinoza) are just returning.

The company security policy mandates that user accounts for employees gone for longer than two weeks be disabled.

In this lab, your task is to:

- Lock the following user accounts:
    - **vedwards**
    - **cflynn**
    - **bkahn**
- Unlock the following user accounts:
    - **mbrown**
    - **bcassini**
    - **aespinoza**
- When you're finished, view the **/etc/shadow** file to verify the changes.

Explanation

Complete this lab as follows:

1. Lock the applicable accounts.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **usermod -L vedwards** or **passwd -l vedwards** and press **Enter** .
    c. Type **usermod -L cflynn** or **passwd -l cflynn** and press **Enter** .
    d. Type **usermod -L bkahn** or **passwd -l bkahn** and press **Enter** .
2. Unlock the applicable accounts.
    a. Type **usermod -U mbrown** or **passwd -u mbrown** and press **Enter** .
    b. Type **usermod -U bcassini** or **passwd -u bcassini** and press **Enter** .
    c. Type **usermod -U aespinoza** or **passwd -u aespinoza** and press **Enter** .
3. Verify your changes by typing **cat /etc/shadow** and pressing **Enter** .
   The inclusion of the exclamation point (**!**) in the password field indicates whether the account is disabled.

4.6.10 Linux User Security and Restrictions (Lesson Video)

Transcript:

FILE NOT FOUND

4.6.11 Configuring Linux User Security and Restrictions (Demo Video)

Transcript:

In this demonstration, we're going to talk about configuring user restrictions. We're going to look at this topic from three different aspects: password aging, setting login limits, and then setting user limits with the ulimit command.

Let's begin by talking about password aging. In today's security environment, you need to be very careful that you configure your passwords to expire after a set period of time. This is called password aging. The key thing to remember here is that the longer a user has the same password, the more it likely it is to be compromised. To prevent this, you need to configure aging for your user password. This is done with the chage command.
Let's look at the 'chage' man page. The syntax for using chage is to enter the chage command followed by a series of options and then the user account that you want to apply those options to. Some of the more useful options you can use are listed right here. First, we have the lowercase -m option, which specifies the minimum number of days between password changes. You also have the uppercase -M option, which specifies the maximum number of days between password changes. And then we have the -w option that specifies the number of warning days a user gets before the password change is required. Go ahead and exit out of the man page.
Let's work through an example. I'm going to enter 'chage' at the shell prompt as my root user. Then I'm going to specify the capital '-M' option to set the maximum number of days between password changes. Let's set that to two months. For many organizations, that would be way too long. A lot of organizations force password changes every 30 days. Let's specify '-w' to specify that a warning is given seven days before the password is about to expire, so the user has plenty of time to make the change before the password actually expires.
Then we must specify who we want to apply this change to. Let's apply it to our rmckay user account. The password is set to expire for the rmckay user in 60 days. Let's use the 'tail' command to verify this by viewing the /etc/shadow file. Look at the very end of the file. Take a look at the rmckay user account. We can see here that the maximum number of days between password changes is 60, and the rmckay user will get seven warning days to change their password before it expires.
The next thing we need to discuss is configuring user limits. Understand that you can configure limits for how many times a user may log in, how much CPU time they use, how much memory they can use, and so on. We configure these limits in the '/etc/security/limits.conf' file. The syntax for this file is shown here. First, we specify the domain. The domain, as noted here, can be a user, or it could be a group.
If you want to use a group, you have to put this '@' sign in front of the group name to indicate that it's a group, or you can use a wild card-- listed here--to apply it to everybody. Scroll down a little bit so we can see more.
Next, we specify the type. As noted here, the type can have two different values. We can either specify a hard limit that can never be exceeded, or we can specify a soft limit that can be exceeded, but just temporarily. Then we specify the item. This is the particular thing that's going to be limited. As you can see down here, we have lots of different options for how we can limit what the item can be.
We can specify core to limit the size of the core dump files. We can specify data to configure a maximum data size and memory for the user's programs. We can specify fsize to restrict the maximum file size. We can specify nofile. We can specify the maximum number of open files. We can specify rss to set the maximum resident set size and memory. We can specify stack to set the maximum stack size. We can specify cpu and set the amount of CPU time that can be used by a single process in minutes. We can specify nproc to specify the number of concurrent processes that will be allowed. Or we can come down here and specify max logins to specify the maximum number of simultaneous logins that we allow for this user. We can also specify priority if we want to, to specify the priority to run user processes with, for this particular user account. Once we specify the item, we scroll back here, and then we specify the value for that item. This is the limit that we're going to configure.
With this in mind, let's go down here, and let's create a new limit. We're going to configure a limit for our rmckay user. I'll enter 'rmckay' for our domain. Next, we want to specify a hard limit, and the item we want to limit is our CPU time. This is the maximum amount of CPU time that a single process run by this user can consume. We must specify how much that is. We do that in minutes. We'll specify 10 minutes. Lets save this by typing "wq!" and hitting Enter. To apply the change, we have to reboot the system.

The last thing we're going to look at the in this demonstration is the ulimit command. To be honest, I don't really care much for the ulimit command. In my experience, it's not as useful as the limits.conf file that we just looked at. The ulimit command does allow you to configure limits on system resources on a per-user basis, much like we did in the limits.conf file. However, be aware that any limits that you configure with ulimit will only affect programs that are launched from within the shell prompt. If, on the other hand, the user was to come over here and launch a graphical application on the desktop, then the limits you specify with ulimit aren't applied. That's why I don't really care for it.
The syntax for ulimit is to enter 'ulimit' followed by the options you want to use and then the limit that you want to specify. There are many different limits you can configure with ulimit. Take a look at the ulimit man page to see a full list of what you can do. For our purposes today, let's just run 'ulimit -a' to, first of all, view the current limits for my user account. You can see what they are right here.
You can also use ulimit to set a limit. Let's set a basic limit. Let's suppose we want to set a soft limit of 100 concurrent processes for my user account. To do this, we would enter' ulimit-s' to set a soft limit. If we wanted to, we could use -h to set a hard limit. Then we have to specify which limit we want to use. We'll use the '-u' option to set the maximum number of processes available to the user. How did I know that that's the right option to use? I looked at the man page. That's what you should do as well. Then we must set that to a specific value. Let's enter '100'. My user account can only have a maximum of 100 concurrent processes. Enter, and the limit is applied.
If we run 'ulimit -a' again, we see that my max user processes have dropped from 7,084 down to 100, which, depending on how the system is being used, would not be anywhere near enough. That's essentially how you use ulimit. Basically, you use ulimit to constrain what users can do so you don't end up with one user hogging all the system resources and not allowing other users on the system to access them.
That's it for this demonstration. In this demo, we talked about how to set up user restrictions on limits. First, we looked at password aging. We looked at login limits. Then we ended this demonstration by looking at the ulimit command.

4.6.12 Configure SELinux (Demo Video)

Transcript:

In this lesson, we're going to review SELinux, or Secure Linux. SELinux allows for a more granular approach to securing Linux from the outside world. This version is installed on most RHEL- or CentOS-based distributions by default. In order to properly configure SELinux, we need to elevate to the root account.

Let's do the sestatus command to see if SELinux is running already. We see the current mode is enforcing. We'll go over that in a little bit.
We use the getenforce command to view the current SELinux mode. This lets us know if the system is enforcing the right properties. We can change enforcement modes with this setenforce command. Type setenforce 0 to change the mode to permissive. This will turn off the protections and controls of SELinux. When we run getenforce again, it's configured to permissive. Type setenforce 1 to change the mode back to enforcing. Now, these changes are temporary. Once we reboot, SELinux will load the configurations in the SELinux config file.
Let's change directories to /etc/selinux and open the config file. We see that SELinux is enforcing, and the Linux type is targeted. This means SELinux will only monitor certain processes, and those processes will be protected. The minimum is a modification of the targeted policy. Once selected, they're protected. And MLS means multi-level security protection. We're just going to look at the targeted type in this demonstration.
Let's exit our text editor. We're going run a program called getsebool. Getsebool is Boolean, meaning that it's either on or off. Now that's add -a for "show all". We see that there are many different Boolean settings for different processes. We're going to look only at the samba processes. Let's run getsebool -a | grep samba to filter only samba processes. The process we're looking to manage here is samba_enable_home_dirs.
So, samba is the process that allows Linux to share files with Windows devices by appearing like a Windows file server. We don't want to enable home directories to be advertised on samba. Let's run setsebool samba_enable_home_dirs=off. Okay, when we rerun the filtered getsebool command, we see that samba_enable_home_dirs is now off.
Again, this is a temporary setting. If we wanted to make the changes permanent, we'd type the same setsebool command but with a dash capital P for "permanent" and then samba_enable_home_dirs=off. Now we press Enter, and when the system reboots, samba home directories will be set to off. Let's move on.
So, files take on different permissions with SELinux. Let's switch to the home directory and use touch file1 to create a file. Let's also see some additional SELinux components. Run ls -lZ file1. We see that the permissions are standard, with Read/write for the user and Read for the group. And then it's Read for everybody else. It's owned by root, and it's owned

by a group called root as well. Now, though, we now have additional information. We have the user type, the role base, and the type of file that it is.
We can change this new information if we want. You might do this because you want to alter who can change the file or something like that. It gives us an additional granularity of control. To do this, you'd use the change context command, called chcon. We're going to change the type from admin_home to a user_home type.
Let's type chcon -t. Now, the -t is what changes the file type, not the role or user type. Okay, continue typing user_home_t file1. We identify the file type and the file we're making the change to. Let's run ls -lZ for file1, and we see that we've changed the type from admin_home to user_home.
Let's say that the changes ended up being a mistake. We could restore the context with a command that changes everything back. We'd run restorecon, specify the file name, and it'd restore the original settings. When we run ls -lZ file1 again, we see the changes.
And that's all for now. In this demonstration, we showed you some SELinux commands, including sestatus, getenforce, and setenforce. Then we modified some Boolean values with getsebool and setsebool. And we also used chcon to change and restore context.

4.6.13 Linux User Security and Restriction Facts

This lesson covers the following topics:

- User security
- User security commands

User Security

When considering user security, keep the following in mind:

- Users should be trained to use secure passwords. Secure passwords use numbers and letters and are more than seven characters in length.
- Passwords should expire periodically but not too often.
- Administrators can limit the resources that the user can access.

User Security Commands

The following table describes Linux commands used to promote user security and restrictions:

Command	Description
chage	Set user passwords to expire. Be aware of the following options: • **-M** sets the maximum number of days before the password expires. • **-W** sets the number of days before the password expires that a warning message displays. • **-m** sets the minimum number of days that must pass after a password has been changed before a user can change the password again.
ulimit	Limits computer resources used for applications launched from the shell. Limits can be hard or soft limits. Soft limits can be temporarily exceeded up to the hard limit setting. Users can modify soft limits, but only the root user can modify hard limits. Options include:

	• **-c** limits the size of a core dump file. The value is in blocks. • **-f** limits the file size of files created using the shell session. The value is in blocks. • **-n** limits the maximum number of files that can be open. • **-t** limits the amount of CPU time a process can use. This is set in seconds. • **-u** limits the number of concurrent processes a user can run. • **-d** limits the maximum amount of memory a process can use. The value is in kilobytes. • **-H** sets a hard resource limit. • **-S** sets a soft resource limit. • **-a** displays current limits. The default shows soft limits.

4.6.14 Practice Questions (Section Quiz)

q_linux_usr_cmds_audit_secp8

You have performed an audit and found an active account for an employee with the username *joer*. This user no longer works for the company.

Which command can you use to disable this account?

Answers:

- ***usermod -L joer**
- usermod -l joer
- usermod -d joer
- usermod -u joer

Explanation:

Use **usermod -L joer** to lock the user's password. Doing so disables the account.

The **usermod -l joer** command changes the account's login name.

The **-d** flag is used for changing the account's home directory.

The **-u** flag is used for changing the account's numeric ID.

q_linux_usr_cmds_login_defs_03_secp8

What information does the /etc/login.defs file contain in a Linux system?

Answers:

- ***User and group ID numbers**
- User account passwords
- List of all installed software
- Network configuration details

Explanation:

User and group ID numbers is the correct answer. The /etc/login.defs file contains values used for the group and user ID numbers. It also includes parameters for password encryption in the shadow file and password expiration values for user accounts.

User account passwords are not stored in the /etc/login.defs file. They are stored in the /etc/shadow file.

The /etc/login.defs file does not contain a list of all installed software. This information can be found in other locations, such as the package manager database.

Network configuration details are not stored in the /etc/login.defs file. These details are typically found in files within the /etc/network directory.

q_linux_usr_cmds_user_01_secp8

One of your users, Karen Scott, has recently married and is now Karen Jones. She has requested that her username be changed from kscott to kjones with no other values changed.

Which of the following commands would accomplish this?

Answers:

- ***usermod -l kjones kscott**
- usermod -l kscott kjones
- usermod -u kjones kscott
- usermod -u kscott kjones

Explanation:

Use the **usermod** command to modify user settings. Use the **-l** flag to signal a change to the username. The correct syntax requires the new username value be given, followed by the old username.

The **-u** flag changes the UID number.

q_linux_usr_cmds_user_02_secp8

An employee named Bob Smith, whose username is bsmith, has left the company. You have been instructed to delete his user account and home directory.

Which of the following commands would produce the required outcome? (Select two.)

Answers:

- ***userdel -r bsmith**
- userdel bsmith
- ***userdel bsmith;rm -rf /home/bsmith**
- userdel -h bsmith
- userdel -Z bsmith

Explanation:

The **userdel -r** command deletes a user's home directory and user account. The **userdel** command by itself does not delete a user's home directory and user account. Executing **rm -rf** on the user's home directory after executing **userdel** removes the home directory.

The **userdel -h** command displays the syntax and options for the **userdel** command.

The **userdel -Z** command removes an SELinux user assigned to the user's login from SELinux login mapping. However, it does not delete the user home directory.

q_linux_usr_cmd_lockout_01_secp8

In the /etc/shadow file, which character in the password field indicates that a standard user account is locked?

Answers:

- !
- !!
- exclamation point
- Exclamation Point
- Exclamation point
- double exclamation point
- Double Exclamation Point

Explanation:

! or *!!* in the password field of /etc/shadow indicates that the account is locked and cannot be used to log in. The /etc/shadow file holds passwords and password expiration information for user accounts.

$ preceding the password identifies the password as an encrypted entry. * indicates a system user account entry (which cannot be used to log in).

q_linux_usr_cmd_lockout_02_secp8

Which of the following utilities could you use to lock a user account? (Select two.)

Answers:

- ***passwd**
- ***usermod**
- useradd
- userdel
- ulimit

Explanation:

Use the following utilities to lock a user account:

- **passwd -l** disables (locks) an account. This command inserts *!!* before the password in the /etc/shadow file.
- **usermod -L** disables (locks) an account. This command inserts *!* before the password in the /etc/shadow file.

The **useradd** command creates new user accounts, and **userdel** deletes user accounts from the system.

The **ulimit** command is used to limit computer resources.

q_linux_usr_cmd_lockout_03_secp8

You suspect that the gshant user account is locked.

Enter the command you would use in a Linux shell to show the status of the user account.

Answers:

- passwd -S gshant
- cat /etc/shadow
- tail /etc/shadow
- more /etc/shadow
- less /etc/shadow

Explanation:

Use **passwd -S gshant** to display the status of the gshant user account.

- LK indicates that the user account is locked.
- PS indicates that the user account has a password.

Viewing the /etc/shadow file also displays whether the user account is disabled. The second field for each entry in the /etc/password file is the password field:

- *$* preceding the password identifies the password as an encrypted entry.
- *!* or *!!* indicates the account is locked and cannot be used to log in.
- * indicates a system account entry, which cannot be used to log in.

q_usradd1_01

What is Paul Denunzio's user ID?

Answers:

- 509
- ***510**
- 517

q_linux_sec_history_secp8

What is the effect of the following command?

chage -M 60 -W 10 jsmith

Answers:

- ***Sets the password for jsmith to expire after 60 days and gives a warning 10 days before expiration.**
- Sets the password for jsmith to expire after 10 days and gives a warning 60 days before expiration.
- Deletes the jsmith user account after 60 days and gives a warning 10 days before expiration.

- Forces jsmith to keep the password for 60 days before changing it while also giving a warning 10 days before expiration.
- Sets the password for jsmith to expire after 60 days and sets a minimum of 10 days before a user can change the password again.

Explanation:

Using **chage -M 60 -W 10 jsmith** sets the password for jsmith to expire after 60 days and gives a warning 10 days before expiration.

Using **chage** sets user passwords to expire. Be aware of the following options:

- **-M** sets the maximum number of days before the password expires.
- **-W** sets the number of days before the password expires that a warning message displays.
- **-m** sets the minimum number of days that must pass after a password change before a user can change the password again.

q_linux_sec_reuse_secp8

Which of the following **chage** option keeps a user from changing their password every two weeks?

Answers:

- ***-m 33**
- -M 33
- -W 33
- -a 33

Explanation:

Using **chage -m 33** forces a user to keep his or her password for 33 days. This sets the minimum number of days that must pass after a password change before a user can change the password again. Be aware of the other **chage** options:

- **-M** sets the maximum number of days before the password expires.
- **-W** sets the number of days before the password expires that a warning message displays.

The **chage -a** option is not a valid option.

q_linux_sec_root_secp8

Which account type in Linux can modify hard limits using the **ulimit** command?

Answers:

- Standard
- Administrator
- ***Root**
- User

Explanation:

Only the root user in Linux can modify hard limits using the **ulimit** command.

Standard and administrator are Windows user types.

Users can modify soft limits but not hard limits using the **ulimit** command.

q_linux_sec_ulimit_01_secp8

Which of the following commands would you use to view the current soft limits on a Linux machine?

Answers:

- ulimit -c
- ulimit -n
- ***ulimit -a**
- ulimit -u

Explanation:

The **ulimit -a** command displays the current limits. The default shows soft limits.

The **ulimit -c** command limits the size of a core dump file.

The **ulimit -n** command limits the maximum number of files that can be open.

The **ulimit -u** command limits the number of concurrent processes a user can run.

q_linux_sec_ulimit_02_secp8

You are a system administrator and you notice that a particular user's processes are consuming an unusually high amount of system resources, causing performance issues for other users.

You decide to use the **ulimit** command to limit the resources available to this user's processes.

Which of the following options would be the MOST effective solution and why?

Answers:

- Use the -n option to limit the maximum number of files that the user can open.
- Use the -u option to limit the number of concurrent processes the user can run.
- Use the -f option to limit the file size of files created using the shell session.
- ***Use the -t option to limit the amount of CPU time a process can use.**

Explanation:

Limiting the amount of CPU time a process can use would be the most effective solution in this case. This would prevent any single process run by the user from consuming too much CPU time and causing performance issues for other users.

Limiting the number of files that the user can open might not be effective if the user's processes are not file-intensive. This option would be more useful if the user was opening a large number of files and causing I/O issues.

Limiting the number of concurrent processes the user can run could be effective if the user is running a large number of processes at the same time. However, if the user is running a small number of resource-intensive processes, this option might not help.

Limiting the file size of files created using the shell session might not be effective if the user's processes are not creating large files. This option would be more useful if the user was creating large files and filling up disk space.

q_linux_sec_user_security_secp8

What should you keep in mind when considering user security in a Linux environment?

Answers:

- ***Users should be trained to use secure passwords.**
- Passwords should never expire.
- Administrators should give users unlimited access to resources.
- Users should be allowed to change their passwords anytime they want without any restrictions.

Explanation:

Users should be trained to use secure passwords is the correct answer. Secure passwords use numbers and letters and are more than seven characters in length. This makes it harder for unauthorized users to guess or crack the passwords.

Passwords should expire periodically but not too often. This is to ensure that even if a password is compromised, it won't be valid for a long time, reducing the risk of unauthorized access.

Administrators can limit the resources that the user can access. This is a principle of least privilege, where users are given only the access they need to perform their tasks. This reduces the risk of accidental or deliberate misuse of resources.

While users should be allowed to change their passwords, there should be a minimum number of days that must pass after a password has been changed before a user can change the password again. This prevents users from quickly cycling through their password history to return to a preferred (possibly weak or compromised) password.

4.7 Linux Groups

As you study this section, answer the following questions:

- Which usermod option changes the secondary group membership?
- Which command removes all secondary group memberships for specific user accounts?
- Which groupmod option changes the name of a group?

In this section, you will learn to:

- Manage Linux groups
- Rename and create groups
- Add users to a group
- Remove a user from a group

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.6 Given a scenario, implement and maintain identity and access management.

Exam	Objective
	• Provisioning/de-provisioning user accounts
TestOut Security Pro	1.1 Manage identity 1.1.2 Manage Linux users and groups

4.7.1 Managing Linux Groups (Demo Video)

Transcript:

For this demonstration, we're going to spend some timeÂ talking about how to manage groups on a Linux system.Â We'll talk about how to add a group,Â modify a group, and remove a group.

Let's begin by adding a new group to the system.Â Let's say we want to add a new groupÂ to the system named Development.Â We do this using the groupadd command. Notice that I'mÂ currently logged in as the rmckay user.Â Standard users on the system aren't allowed to add other users or groups. You can only do that as a user with sudo privileges or as the root user. For this demo, we're going to use â€˜su â€“ root'. You'll be prompted for the root password. Click Enter.
Before we add a new group to the system, I need to pointÂ out that there are several default parameters that areÂ automatically assigned to a new group when it's created.Â These defaults are stored in the /etc/login.defs file. Let's type â€˜less /etc/login.defs' to take a look.
First of all, the group ID number that's automaticallyÂ assigned when the group is created is defined here.Â They'll begin at 1000 on this system, so theÂ first group created will have a group ID of 1000.Â The next one will be 1001, the next one 1002, and so on.Â Likewise, if we created a system group on theÂ system, the group ID numbering would begin at 201. Every group ID is unique to the group.
let's use 'groupadd' and then specifyÂ the name of the group, 'Development'.Â We'll check to see if our group was created by typing â€˜tail /etc/group'.Â The X here tells us that it's using the gshadow passwordÂ file, but we haven't assigned any passwords.Â Here's the group ID number that's beenÂ assigned to that group. But notice, over here,Â that there are no members of that group yet.
In order to add members to a group, you mustÂ modify the group, but the actual command you use willÂ vary based upon what distribution you're using.
We used a capitalÂ D in the name of our group when it should have been a lowercase d like the rest of the groups.Â Let's make it match by modifyingÂ the group name: 'groupmod -n'.Â Now we specify the new group name that we want to use,Â 'development', and then we specify the name of the existingÂ group that we want to modify, 'Development', and click Enter.Â Let's view the end of the group file with a â€˜tail /etc/group' command. We see that the group has been renamed with a lowercase d.Â We know that's the same group and notÂ a new one because the group ID number,Â which uniquely identifies that group, is the same number as before.
To add usersÂ to a group, you can use the usermod command. Instead of modifying the group and adding usersÂ to it, we're going to modify a userÂ and tell it that it's now a member of this group.Â Let's go ahead and view the man page for usermod by typing â€˜man usermod'.Â As we scroll down, we should see the -gÂ options, lowercase g and uppercase G.
The lowercase g option is used toÂ modify the user's primary group, the default group.Â Remember that every user account on the Linux systemÂ has one primary group associated with it, and only one.Â If you want to change the default group for a userÂ account, you use the lowercase g option.Â Also remember that a Linux user can be a member of manyÂ other groups as well; these are our supplementary groups.Â If we want to add a user as a member of a supplementaryÂ group, we can use the uppercase G option instead.
Here's a very important thing that you have toÂ remember (and it trips up a lot of new LinuxÂ administrators): if you specify the usermod commandÂ and use the uppercase G option followed by the name ofÂ a group, whatever you specify with the G option willÂ overwrite whatever group memberships that user already has.Â If I have a user that's a

member of three differentÂ groups already, and I use the uppercase G optionÂ with the usermod command and specify an additional groupÂ that I want to make the user a member of, what actuallyÂ happens is the existing group memberships are removed andÂ replaced with the one group membership that I specify.Â If you want to add an additional group membership andÂ not replace an existing group membership, you need toÂ use the uppercase G option along with the -a option.Â The -a option appends the new group to theÂ list of group memberships instead of replacing it.
Let's go ahead and add development as aÂ supplementary group to my rmckay user account.Â To do this, I type 'usermod -G' followed by theÂ name of the group that I want to add, 'development'.Â Then I specify '-a' to indicate that I don't want toÂ remove any other group memberships already in place--I just want to add this as a new group membership.Â And then the name of the user account thatÂ I want to add the membership to, 'rmckay'. If we type â€˜groups rmckay', we can tell that the rmckay user is a part of the development group and hasn't lost any previously added groups.
Up to this point, we've talked about how to add a newÂ group, and we've talked about how to modify a group.Â Let's end this demonstration by discussingÂ how to remove a group from the system.Â This is done using the groupdel command. All I have to do is type 'groupdel development'.Â If we tail the group file with a â€˜tail /etc/group' command, we seeÂ that the development group is gone.
That's it for this demonstration.Â In this demo, we talked about how toÂ manage Linux groups from the command line.Â We talked about how to add aÂ group with the groupadd command.Â We talked about how to modify a groupÂ with the groupmod and usermod commands.Â Then we ended this demonstration by talking aboutÂ how to remove a group with the groupdel command.

4.7.2 Linux Group Commands

This lesson covers the following topic:

- Linux group commands

Linux Group Commands

Use the following commands to manage group accounts and group membership:

Command	Function
groupadd	Creates a new group. The following options override the settings found in the /etc/login.defs file: • **g** defines the group ID (GID). • **p** defines the group password. • **- r** creates a system group.
groupmod	Modifies the existing group.
groupdel	Modifies the system account files by deleting all entries that refer to the specified group. The named group must exist. You cannot remove the primary group of any existing user. You must remove the user before you remove the group.
gpasswd	Changes a group password. • **groupname** prompts for a new password. • **- r** removes a group password.
newgrp	Is used to change the current group ID during a login session. If the optional - flag is given, the user's environment will be reinitialized as though the user had logged in. Otherwise, the current environment,

Command	Function
	including the working directory, remains unchanged. You can use this when working in a directory where all the files must have the same group ownership.
usermod	Modifies group membership for the user account. Be aware of the following options: • **g** assigns a user to a primary group. • **G** assigns a user to a secondary group (or groups). Follow the command with a comma-separated list of groups. • **aG** assigns a user to a secondary group (or groups) by appending the group to any group the user already belongs to. Follow the command with a comma-separated list of groups. • **- G** "" removes the user from all secondary group memberships. Do not include a space between the quotes.
groups	Display the primary and secondary group membership for the specified user account.

4.7.3 Rename and Create Groups (Simulation)

Scenario

Currently, all the salespeople in your company belong to a group called sales. The VP of sales wants two sales groups, a western sales division, and an eastern sales division.

In this lab, your task is to:

- Rename the sales group to **western_sales_division** .
- Create the **eastern_sales_division** group.
- Remove **aespinoza** as a member of the **western_sales_division** group.
- Assign **aespinoza** as a member of the **eastern_sales_division** group.
- When finished, view the **/etc/group** file or use the **groups** command to verify the changes.

Explanation

Complete this lab as follows:

1. Rename the sales group *western_sales_division* and create the *eastern_sales_division* group.
   a. From the Favorites bar, select **Terminal** .
   b. At the prompt, type **groupmod -n western_sales_division sales** and press **Enter** .
   c. Type **groupadd eastern_sales_division** and press **Enter** .
2. Modify the group membership as needed.
   - Type **usermod -G eastern_sales_division aespinoza** and press **Enter** .

     When you assign aespinoza to the *eastern_sales_division* group using the **usermod -G** option, the user account is removed from the *western_sales_division* group.
3. Use **cat /etc/group** or **groups aespinoza** to verify aespinoza's group membership.

4.7.4 Add Users to a Group (Simulation)

Scenario

Maggie Brown (mbrown) and Corey Flynn (cflynn) have recently been hired in the human resources department. You have already created their user accounts.

In this lab, your task is to:

- Add the **hr** group as a secondary group for the mbrown and cflynn user accounts.
- When finished, view the **/etc/group** file or use the **groups** command to verify the changes.

When the **-g** switch is used with the **usermod** command, it sets the primary group membership, not the secondary one.

Explanation

Complete this lab as follows:

1. Add users to the hr group.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **usermod -G hr mbrown** and press **Enter** .
    c. Use **usermod -G hr cflynn** and press **Enter** .
2. Verify the group membership for the users added to each group.
    a. Use **groups mbrown** and press **Enter** .
    b. Use **groups cflynn** and press **Enter** .

4.7.5 Remove a User from a Group (Simulation)

Scenario

Corey Flynn (cflynn) currently belongs to several groups. Due to some recent restructuring, he no longer needs to be a member of the hr group.

To preserve existing group membership, use the **usermod -G** command to list all groups to which the user must belong. Do not include the primary group name in the list of groups.

In this lab, your task is to:

- Remove **cflynn** from the hr group.
- Preserve all other group memberships.
- View the **/etc/group** file or use the **groups** command to verify the changes.

Explanation

Complete this lab as follows:

1. View a list of all groups to which Cory Flynn belongs.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **groups cflynn** and press **Enter** .
    Notice that cflynn currently belongs to the *mgmt1* , *hr* , and *it* secondary groups. The cflynn group is the user's primary group.
2. Change and verify Cory Flynn's group membership.
    a. Type **usermod -G mgmt1,it cflynn** and press **Enter** .

b. Type **groups cflynn** and press **Enter** .
Cory now only belongs to the *mgmt1* and *it* groups.

4.7.6 Practice Questions (Section Quiz)

q_linux_grps_cmds_add_secp8

You are the administrator for a small company, and you need to add a standard new group of users to the system. The group's name is **sales**.

Which command accomplishes this task?

Answers:

- addgroup sales
- groupadd -r sales
- ***groupadd sales**
- addgroup --system sales

Explanation:

Use the **groupadd** utility to add a group to the system. By default, the group is added with an incrementing number above those reserved for system accounts.

If you use the **-r** option, the account is added as a system account (with a reserved group id number). Because this is a group that is created for users, the **-r** option should not be used.

The **--system** option adds a system group. However, in this scenario, you need to create a normal user group, not a system group.

q_linux_grps_cmds_del_secp8

You have a group named **temp_sales** on your system. The group is no longer needed, so you should remove it.

Which of the following commands should you use?

Answers:

- ***groupdel temp_sales**
- groupmod -R temp_sales
- groupmod -n temp_sales
- newgroup -R temp_sales

Explanation:

Use **groupdel** to delete a group from a Linux system.

The **newgroup** command logs the user into a group with the group password, but this command does not contain a **-R** option. The **groupmod** command modifies the existing group.

Be aware of the following options:

- **-A** adds specified users to the group (SUSE distribution).
- **-R** removes specified users from the group (SUSE distribution).
- **-n** changes the name of a group.

q_linux_grps_cmds_groupadd_secp8

Which of the following commands creates a new group and defines the group password?

Answers:

- groupadd -g
- ***groupadd -p**
- groupadd -r
- groupadd -c

Explanation:

The **groupadd -p** command creates a new group while defining the group password.

The **groupadd -g** command creates a new group while defining the GUID.

The **groupadd -r** command creates a new system group.

The **groupadd -c** command is not a valid command.

q_linux_grps_cmds_groups_secp8

You want to see which primary and secondary groups the dredford user belongs to. Enter the command you would use to display group memberships for **dredford**.

Answers:

- groups dredford
- groups dredford

Explanation:

To display the primary and secondary group membership for a specified user account, use the **groups** command. In this case, you would enter:

groups dredford

q_linux_grps_cmds_login_secp8

Using the **groupadd -p** command overrides the settings found in which file?

Answers:

- /etc/logins.txt
- /root/logins.defs
- /usr/logins.txt

- ***/etc/login.defs**

Explanation:

Using the **groupadd** command with the **-p** option overrides the default settings found in the /etc/login.defs file. The file is not located in the **/root/** directory.

There is no **logins.txt** file in Linux.

q_linux_grps_cmds_newgrp_secp8

Which of the following commands is used to change the current group ID during a login session?

Answers:

- ***newgrp**
- usermod
- groups
- groupmod

Explanation:

The **newgrp** command is used to change the current group ID during a login session. If the optional - flag is given, the user's environment is reinitialized as though the user had logged in. Otherwise, the current environment (including the current working directory) remains unchanged. You can use this when working in a directory in which all the files must have the same group ownership.

The **usermod** command modifies group membership for a user account.

The **groups** command displays the primary and secondary group membership for the specified user account.

The **groupmod** command modifies the existing group.

q_linux_grps_cmds_pass_secp8

You have a group named Research on your system that needs a new password because a member of the group has left the company.

Which of the following commands should you use?

Answers:

- ***gpasswd Research**
- newpasswd Research
- gpasswd research
- groupmod -p Research

Explanation:

Use **gpasswd Research** to be prompted to enter a new password for the Research group.

Group names are case-sensitive, so **gpasswd research** won't change the password for the Research group.

The **groupmod** command does not have a switch that can be used to change passwords.

The **newpasswd** option is not a valid Linux command.

q_linux_grps_cmds_primary_secp8

You are attempting to delete the temp group but are unable to.

Which of the following is the MOST likely cause?

Answers:

- All users have already been deleted.
- Groups cannot be deleted.
- ***The primary group of an existing user cannot be deleted.**
- The secondary group of an existing user cannot be deleted.

Explanation:

You cannot remove the primary group of any existing user. You must remove the user before you remove the group.

Deleting all users would not prevent a group from being deleted.

Groups can be deleted using the **groupdel** command.

Secondary groups of a user can be deleted. This event would not prevent a group from bring deleted.

q_linux_grps_cmds_remove_secp8

Which of the following commands removes a user from all secondary group memberships?

Answers:

- usermod -g
- usermod -G
- ***usermod -G ""**
- usermod -aG

Explanation:

usermod - G "" removes the user from all secondary group memberships. Do not include a space between the quotes.

usermod -g assigns a user to a primary group.

usermod -G assigns a user to a secondary group.

usermod -aG assigns a user to a secondary group (or groups) by appending the group to any which the user already belongs to. Follow the command with a comma-separated list of groups.

q_linux_grps_cmds_usermod_secp8

Which of the following commands assigns a user to a primary group?

Answers:

- ***usermod -g**
- usermod -G
- groupadd -g
- groupadd - r

Explanation:

The **usermod -g** command assigns a user to a primary group.

The **usermod -G** command assigns a user to a secondary group.

The **groupadd -g** command creates a new group while defining the GUID.

The **groupadd -r** command creates a new system group.

4.8 Remote Access

As you study this section, answer the following questions:

- How does EAP differ from CHAP?
- How can remote access and tunneling be secured?
- What is the difference between RADIUS and TACACS+?

In this section, you will learn to:

- Configure a RADIUS solution

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	3.2 Given a scenario, apply security principles to secure enterprise infrastructure.
	• Secure communication/access
	  ○ Virtual private network (VPN)
	  ○ Remote access
	  ○ Tunneling
	  ▪ Internet Protocol Security (IPSec)
	
	4.1 Given a scenario, apply common security techniques to computing resources.
	• Wireless security settings
	  ○ AAA/Remote Authentication Dial-In User Service (RADIUS)

Exam	Objective
TestOut Security Pro	2.2 Harden network devices • 2.2.3 Configure and access a virtual private network (VPN)

4.8.1 Remote Access (Lesson Video)

Transcript:

In this lesson we're going to talk about remote access. Remote access involves the infrastructure, protocols, and software that allow a host to join a local network from a physically remote location or to establish a session on a host over a network.

Remote access networking means the user's device does not make a direct cable or wireless connection to the network. The connection occurs over or through an intermediate network.
Historically, remote access used analog modems connecting over the telephone system, but these days, most remote access is implemented as a virtual private network , or VPN, running over an ISP's network.
This client-to-site VPN topology is the "telecommuter" model, allowing homeworkers and employees working in the field to connect to the corporate network. The VPN protocol establishes a secure tunnel to keep the contents private, even when the packets pass over ISPs' routers.
The VPN client host first connects to a VPN gateway using an internet connection. The VPN gateway then authenticates the user and creates a secure encrypted tunnel. The VPN client traffic is routed over the network and can access authorized services.
A VPN can also be deployed in a site-to-site model to connect two or more private networks. While remote access VPN connections are typically initiated by the client, a site-to-site VPN is configured to operate automatically. The gateways exchange security information using whichever protocol the VPN is based on. This establishes a trust relationship between the gateways and sets up a secure connection through which to tunnel data. Hosts at each site do not need to be configured with any information about the VPN. The routing infrastructure at each site determines whether to deliver traffic locally or send it over the VPN tunnel.
Remote desktop refers to a technology that allows users to access and control a computer or device remotely. It enables individuals to connect to another computer over a network, typically the internet, and operate it as if they were physically in front of it.
With a remote desktop, you can access files, programs, and resources on a remote computer without having to be physically present at that location. This technology benefits remote work, technical support, collaboration, and access to personal or work-related resources from anywhere.
A remote access VPN joins the user's PC or smartphone to a remote private network via a secure tunnel over a public network. Remote access can also be a means of connecting to a specific computer over a network. This type of remote access involves connecting to a terminal server on a host using software that transfers shell data only. The connection could be a client and terminal server on the same local network or across remote networks.
A graphical remote access tool sends screen and audio data from the remote host to the client. It transfers mouse and keyboard input from the client to the remote host. Microsoft's Remote Desktop Protocol (RDP) can access a physical machine on a one-to-one basis.
That's it for this lesson. In this lesson, we discussed remote access architecture and remote desktop and how they help widen the geographical reach of an individual or an organization without the need for travel.

4.8.2 Remote Access Facts

Remote access involves the infrastructure, protocols, and software that allow a host to join a local network from a physically remote location or to establish a session on a host over a network.

This lesson covers the following topics:

- Remote access architecture
- Remote desktop

Remote Access Architecture

Remote access networking means that the user's device does not make a direct cable or wireless connection to the network. The connection occurs over or through an intermediate network.

Historically, remote access used analog modems connecting over the telephone system. These days, most remote access is implemented as a virtual private network (VPN) , running over Internet Service Provider (ISP) networks.

With a remote access VPN, clients connect to a VPN gateway on the edge of the private network. This **client-to-site** VPN topology is the "telecommuter" model, allowing homeworkers and employees working in the field to connect to the corporate network. The VPN protocol establishes a secure tunnel to keep the contents private, even when the packets pass over ISPs' routers.

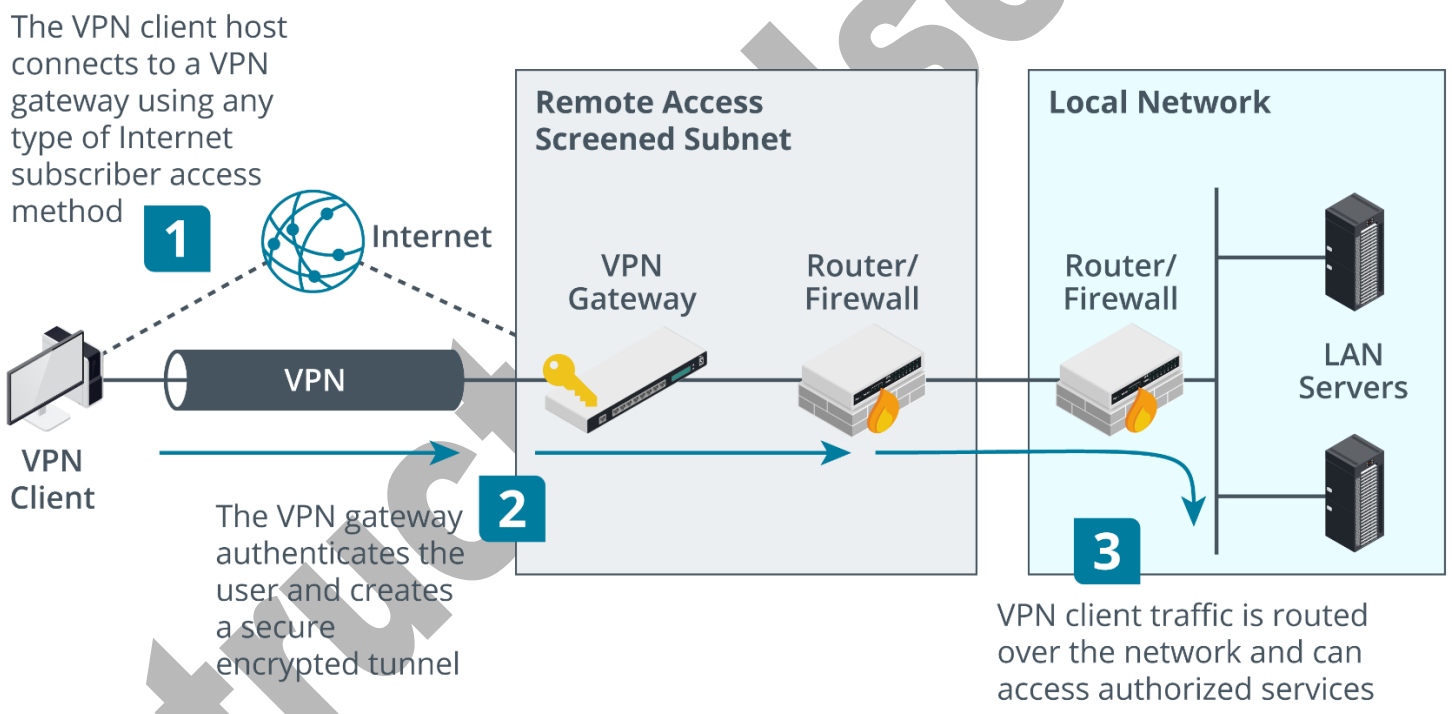

Remote access VPN. (Images © 123RF.com.)

A VPN can also be deployed in a **site-to-site** model to connect two or more private networks. Whereas remote access VPN connections are typically initiated by the client, a site-to-site VPN is configured to operate automatically. The gateways exchange security information using whichever protocol the VPN is based on. This establishes a trust relationship between the gateways and sets up a secure connection through which to tunnel data. Hosts at each site do not need to be configured with any information about the VPN. The routing infrastructure at each site determines whether to deliver traffic locally or send it over the VPN tunnel.

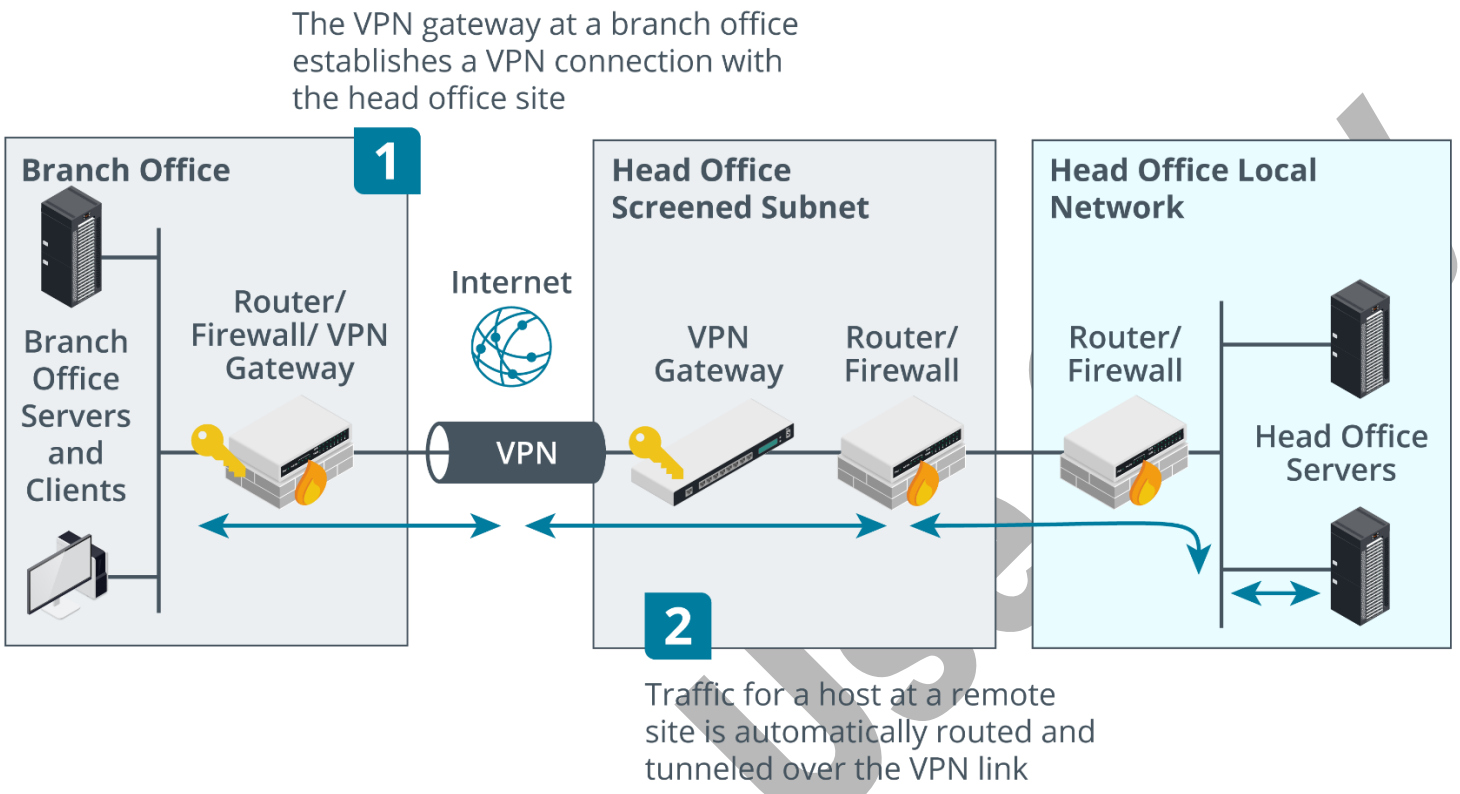

Site-to-site VPN. (Images © 123RF.com.)

A third topology is a **host-to-host tunnel** . This is a means of securing traffic between two computers where the private network is not trusted.

Several VPN protocols have been used over the years. Legacy protocols, such as the Point-to-Point Tunneling Protocol (PPTP), have been deprecated because they do not offer adequate security. Transport Layer Security (TLS) and Internet Protocol Security (IPsec) are now the preferred options for configuring VPN access.

Remote Desktop

A remote access VPN joins the user's PC or smartphone to a remote private network via a secure tunnel over a public network. Remote access can also be a means of connecting to a specific computer over a network. This type of remote access involves connecting to a terminal server on a host using software that transfers shell data only. The connection could be a client and terminal server on the same local network or across remote networks.

A graphical remote access tool sends screen and audio data from the remote host to the client and transfers mouse and keyboard input from the client to the remote host. Microsoft's Remote Desktop Protocol (RDP) can be used to access a physical machine on a one-to-one basis.

Alternatively, a site can operate a remote desktop gateway that facilitates access to virtual desktops or individual apps running on the network servers. RDP connections are encrypted by default. There are several popular alternatives to Remote Desktops. Most support remote access to platforms other than Windows (macOS and iOS, Linux, Chrome OS, and Android, for instance). Examples include TeamViewer (teamviewer.com/en) and Virtual Network Computing (VNC) , which is implemented by several different providers (notably realvnc.com/en).

In the past, these remote desktop products required a dedicated client app. Remote desktop access can now just use a web browser client. The canvas element introduced in HTML5 allows a browser to draw and update a desktop with relatively little lag. It can also handle audio. This is referred to as an HTML5 VPN or as a clientless remote desktop gateway (guacamole.apache.org). This solution uses a protocol called WebSocket, which enables bidirectional messages to be sent between the server and client without requiring the overhead of separate HTTP requests.

4.8.3 Configuring a RADIUS Solution (Demo Video)

Transcript:

Now, let's take a look at creating RADIUS clients. I'm going to go into my Network Policy Server management console. You can see up here that we have a list of RADIUS Clients and Remote RADIUS Servers. Whoever's talking to RADIUS needs to be configured as a RADIUS client. The only exception is if RADIUS is installed on the same server as Remote Access, which is actually the case in my environment. In that case, I don't need to configure Routing and Remote Access as a RADIUS client because it's on the same server.

When you install the Network Policy and Access server role, it changes what you see in Routing and Remote Access. Let's go into the Properties of my server here and to Security. Before I installed NPS, I had some drop-down combo boxes here that said, "Do you want to use Windows authentication or RADIUS authentication?" Now that I've installed NPS, I have to use it to configure authentication and accounting providers.
The most I can do here is change my authentication method. Installing NPS changes the interface for Routing and Remote Access. Let's say I had something else that was going to talk to NPS, a wireless router or wireless switch, a wired switch, or some other server that's running Routing and Remote Access, and it wants to talk to this NPS. I'd need to define them as a RADIUS client. I can define RADIUS client templates or right-click here and do a New RADIUS client. I can base this RADIUS client on an existing template. In that case, it will just fill out the information for me, or I can put in the information manually.
Then, I need to provide a Shared secret. This will be put in at both the RADIUS client and here at NPS, where I define the RADIUS client. It's just a phrase used to encrypt the connection. I can base it on templates or manually type in the Shared secret. No matter what, I have to define the RADIUS client for it to talk to NPS.
After this, you need to go to the RADIUS client itself—the switch, the other router, the other server, whatever's the RADIUS client—and tell it to talk to NPS on this server. How you do that depends on what's using this NPS for authentication, authorization, and accounting. It's always two steps: defining the RADIUS client on NPS and then, at the RADIUS client, defining RADIUS and pointing it to the NPS server.
There are several authentication methods using certificates for NPS that will need to be set up in order to have a secure authentication method, such as MS-CHAP V2, MS-CHAP, or CHAP.
That's it for this demo. In this demo, we went over the basics of configuring a RADIUS client and looked at how RADIUS affects routing and remote access.

4.8.4 RADIUS and TACACS+ Facts

This lesson covers the following topics:

- AAA server
- RADIUS
- TACACS+

AAA Server

An AAA server handles user requests for access to computer resources. A remote access server typically controls client access to remote systems. Clients might be restricted to accessing resources only on the remote access server, or they might be allowed to access resources on other hosts on the private network. Two standard AAA server solutions include RADIUS and TACACS+.

Remote access policies identify authorized users and other required connection parameters. In a small implementation, you typically define user accounts and remote access policies on the remote access server. With this configuration, you must define user accounts and policies on each remote access server. For larger deployments with multiple remote access servers, you can centralize the administration of remote access policies using an AAA server.

The remote access server receives connection requests from remote clients. It forwards them to the AAA server for approval or denial. The policies you define on the AAA server apply to all clients connected to all remote access servers.

RADIUS

Microsoft servers use RADIUS for centralized remote access administration. When using RADIUS, be aware that RADIUS:

- Combines authentication, authorization, and accounting. All three must be implemented through the RADIUS system.
- Allows for the separation of accounting to different servers. However, authentication and authorization remain combined on a single server.
- Supports PPP, CHAP, and PAP.
- Uses a challenge-response method for authentication.
- Does not transmit passwords in cleartext between the RADIUS client and the RADIUS server.
- A shared secret is used between the RADIUS server and the RADIUS client.
- The password is hashed, and the hash is added before it is transmitted.
- Encrypts only the password using MD5.
- Uses UDP ports 1812 and 1813 and can be vulnerable to buffer overflow attacks.
- Often uses vendor-specific extensions. RADIUS solutions from different vendors might not be compatible.

For example, a NAS device (RADIUS client) is configured with the IP address of the RADIUS server and a shared secret. This allows the client to authenticate to the server. Remember that the client is the access device (switch, access point, or VPN gateway), not the user's PC or laptop. A generic RADIUS authentication workflow proceeds as follows:

- The user's device (the supplicant) connects to the NAS appliance, such as an access point, switch, or remote access server.
- The NAS prompts the user for their authentication credentials. RADIUS supports PAP, CHAP, and EAP. Most implementations now use EAP, as PAP and CHAP are not secure. If EAP credentials are required, the NAS enables the supplicant to transmit EAP over LAN (EAPoL) data but not any other type of network traffic.
- The supplicant submits the credentials as EAPoL data. The RADIUS client uses this information to create an Access-Request RADIUS packet, encrypted using the shared secret. It sends the Access-Request to the AAA server using UDP on port 1812 (by default).
- The AAA server decrypts the Access-Request using the shared secret. If the Access-Request cannot be decrypted (because the shared secret is not correctly configured, for instance), the server does not respond.
- With EAP, Access-Challenge and Access-Request packets will be exchanged as the authentication method is set up and the credentials verified. The NAS acts as a pass-thru, taking RADIUS messages from the server and encapsulating them as EAPoL to transmit to the supplicant.
- At the end of this exchange, if the supplicant is authenticated, the AAA server responds with an Access-Accept packet; otherwise, an Access-Reject packet is returned.

Optionally, the NAS can use RADIUS for accounting (logging). Accounting uses port 1813. The accounting server can be different from the authentication server.

TACACS+

TACACS+ was originally developed by Cisco for centralized remote access administration. TACACS+:

- Provides three protocols, one each for authentication, authorization, and accounting. This allows each service to be provided by a different server.
- Uses TCP port 49.
- Encrypts the entire packet contents, not just authentication packets. The client-server dialogs are also encrypted.
- Supports more protocol suites than RADIUS.
- Requires remote access servers to become TACACS+ clients to the backend TACACS+ server, like a RADIUS solution.

TACACS was initially developed in 1984 by BBN Technologies. The protocol standard, TACACS+, was developed by Cisco Systems but is supported by many vendors, such as BlueCat Networks, IBM, Netgear, and more.

TACACS and Extended Terminal Access Controller Access-Control System (XTACACS) are older protocols developed before TACACS+. While they sound similar, they are different and less secure protocols.

4.8.5 Practice Questions (Section Quiz)

q_remote_acc_client_to_site_secp8

A company is planning to implement a remote access architecture to allow its employees to work from home.

The company has a central office where all its servers and applications are located. The employees need to access these resources securely from their home computers.

Which remote access architecture would be the most suitable for this scenario?

Answers:

- Site-to-site VPN topology
- Host-to-host tunnel topology
- ***Client-to-site VPN technology**
- Transport Layer Security (TLS)

Explanation:

Client-to-site VPN technology is the correct answer. Client-to-site VPN technology allows individual users to securely connect to a private network from a remote location. It establishes a secure tunnel between the user's device (the client) and the network's gateway (the site). This enables users to access resources, applications, and services within the private network as if they were physically present at the site. Client-to-site VPNs are commonly used for remote work.

A site-to-site VPN topology is used to connect two or more geographically separate networks over a public network such as the internet. This is typically used by businesses with multiple branch offices or organizations that must securely connect with external partners, not individual users.

A host-to-host tunnel topology is used for secure communication between specific hosts or endpoints. This type of VPN is typically used for secure file transfers, remote access to specific systems, or for establishing secure communication channels between devices in different locations. It is not suitable for connecting individual users to a central network.

While TLS is a protocol that provides privacy and data integrity between two communicating applications, it is not a remote access architecture. TLS is used within many VPN technologies for encryption and integrity checks, but it does not provide the VPN functionality itself.

q_remote_acc_host_to_host_secp8

A tech company is developing a new software product. The development team is distributed across different locations and needs to securely access and work on specific systems located in the company's main office.

The team members need to establish secure communication channels between their individual devices and the specific systems in the office.

Which remote access architecture would be the most suitable for this scenario?

Answers:

- Client-to-site VPN technology
- Site-to-site VPN topology
- ***Host-to-host tunnel topology**
- Virtual network computing (VNC)

Explanation:

Host-to-host tunnel topology is the correct answer. In a host-to-host tunnel topology, individual devices establish a secure tunnel between each other over a public network. This type of VPN is typically used for secure communication between specific hosts or endpoints, which is exactly what the team members need in this scenario.

Client-to-site VPN technology is typically used when individual users need to securely connect to a private network from a remote location. In this scenario, the team members need to establish secure communication channels with specific systems, not the entire network.

A site-to-site VPN topology is used to connect two or more geographically separate networks over a public network such as the Internet. This is typically used by businesses with multiple branch offices or organizations that must securely connect with external partners, not for secure communication between specific hosts or endpoints.

While VNC allows for remote control of another computer, it does not inherently provide the secure, encrypted tunnels needed for this scenario. VNC is a type of remote desktop software, not a remote access architecture. It would allow the team members to control the specific systems, but it would not provide the secure communication channels they need.

q_remote_acc_rdp_secp8

A multinational corporation wants to enable its IT support team to provide remote assistance to employees across various locations. The support team needs to be able to take control of the employees' computers to troubleshoot and resolve issues.

The corporation primarily uses Windows-based systems.

Which technology would be the MOST suitable for this purpose?

Answers:

- Simple Network Management Protocol (SNMP)
- ***Remote Desktop Protocol (RDP)**
- Transport Layer Security (TLS)
- Remote Authentication Dial-in User Service (RADIUS)

Explanation:

Remote Desktop Protocol (RDP) is the correct answer. RDP is a proprietary protocol developed by Microsoft that allows a user to connect to another computer over a network connection in a graphical interface. This makes it ideal for IT support teams to remotely control and troubleshoot issues on employees' computers.

SNMP is an application protocol for monitoring and managing network devices. While it can provide valuable information about network devices, it does not allow for remote control of a user's desktop.

TLS is a cryptographic protocol that provides secure communication over a computer network. While it is used to secure communications, it does not provide the functionality to remotely control a user's desktop.

RADIUS is a networking protocol that provides centralized authentication, authorization, and accounting (AAA) management for users who connect and use a network service. While it is used for managing remote and wireless authentication infrastructures, it does not allow for remote control of a user's desktop.

q_remote_acc_remote_secp8

You often travel away from the office. While traveling, you would like to use your laptop computer to connect directly to a server in your office and access files. You want the connection to be as secure as possible.

Which type of connection do you need?

Answers:

- ***Remote access**
- Internet
- Virtual private network
- Intranet

Explanation:

Use a remote access connection to connect directly to a server at a remote location.

You could use a virtual private network (VPN) connection through the internet to connect to the server security. However, the connection would involve connecting first to the internet through a local ISP and then establishing a VPN connection to the server. While the VPN connection through the internet is secure, it is not as secure as a direct remote connection to the server.

An intranet is an internal network that only internal users can access.

q_remote_acc_tls_secp8

Which of the following protocols is primarily used for secure remote access to a network by creating an encrypted tunnel over the internet?

Answers:

- Simple Network Management Protocol (SNMP)
- Remote Desktop Protocol (RDP)
- ***Transport Layer Security (TLS)**
- Hypertext Transfer Protocol (HTTP)

Explanation:

Transport Layer Security (TLS) is the correct answer. TLS is a protocol that provides privacy and data integrity between two communicating applications. It's used to create an encrypted tunnel for secure remote access to a network over the internet.

SNMP is primarily used for managing and monitoring network devices, not for secure remote access to a network.

Remote Desktop Protocol (RDP) is primarily used for remote desktop access, allowing a user to control another computer over a network connection. While it does provide a form of remote access, it is not primarily used for creating an encrypted tunnel for secure network access.

Hypertext Transfer Protocol (HTTP) is used for transmitting hypertext over the internet, not for secure remote access to a network. HTTP does not provide the necessary encryption for secure remote access.

q_remote_acc_vpn_sol_secp8

A global pharmaceutical company's IT team needs a secure solution for remote employees to access internal company resources from home.

The solution must require user authentication, encapsulate and encrypt all traffic between the user and the internal network, and establish a secure tunnel.

Which solution should the team choose?

Answers:

- ***Virtual Private Network (VPN)**
- Secure Shell (SSH)
- Simple Network Management Protocol (SNMP)
- Remote Desktop Protocol (RDP)

Explanation:

A VPN creates a secure private connection between the remote user's device and the company's internal network that requires authentication and uses a network tunnel.

SSH offers secure, encrypted command-line access and file transfer but does not encapsulate all traffic and does not create a network tunnel.

SNMP is a device management protocol for IP networks, not for creating a secure connection for a remote user.

RDP allows a user to remotely access and control another machine but does not encrypt all traffic or create a network tunnel. RDP focuses on direct connection to a single machine.

q_radius_tacacs_aaa_secp8

What is the primary function of an AAA server in a network?

Answers:

- To provide email services to all users
- ***To handle user requests for access to computer resources**
- To store all data and files in the network
- To provide internet connectivity to all devices in the network

Explanation:

To handle user requests for access to computer resources is the correct answer. An AAA server (authentication, authorization, and accounting) handles user requests for access to computer resources. It typically controls client access to remote systems and can centralize the administration of remote access policies.

An AAA server does not provide email services. Email services are typically provided by an email server.

An AAA server does not store all data and files in the network. This is typically the role of a file server or a database server.

An AAA server does not provide internet connectivity. This is typically the role of a router or a modem.

q_radius_tacacs_radius_secp8

RADIUS is primarily used for what purpose?

Answers:

- Managing access to a network over a VPN
- ***Authenticating remote clients before access to the network is granted**
- Managing RAID fault-tolerant drive configurations
- Controlling entry-gate access using proximity sensors

Explanation:

Remote Authentication Dial-In User Service (RADIUS) is primarily used for authenticating remote clients before access to a network is granted.

RADIUS is based on RFC 2865 and maintains client profiles in a centralized database.

RADIUS offloads the authentication burden for dial-in users from the normal authentication of local network clients.

For environments with a large number of dial-in clients, RADIUS provides improved security, easier administration, improved logging, and alleviated performance impact on LAN security systems.

q_radius_tacacs_tac_char_01_secp8

Which of the following are characteristics of TACACS+? (Select two.)

Answers:

- ***Uses TCP**
- Uses UDP
- ***Allows three different servers (one each for authentication, authorization, and accounting)**
- Allows two different servers (one for authentication and authorization and another for accounting)
- Can be vulnerable to buffer overflow attacks

Explanation:

TACACS+ was originally developed by Cisco for centralized remote access administration. TACACS+:

- Provides three protocols (one each for authentication, authorization, and accounting). This allows each service to be provided by a different server.
- Uses TCP.
- Encrypts the entire packet contents.
- Supports more protocol suites than RADIUS.

RADIUS is used by Microsoft servers for centralized remote access administration. RADIUS:

- Combines authentication and authorization using policies to grant access.
- Uses UDP.
- Encrypts only the password.
- Often uses vendor-specific extensions. RADIUS solutions from different vendors might not be compatible.
- Uses UDP ports 1812 and 1813 and can be vulnerable to buffer overflow attacks.

q_radius_tacacs_tac_char_02_secp8

Which of the following is a characteristic of TACACS+?

Answers:

- ***Encrypts the entire packet, not just authentication packets**
- Requires that authentication and authorization are combined in a single server
- Uses UDP ports 1812 and 1813
- Supports only TCP/IP

Explanation:

TACACS+ was originally developed by Cisco for centralized remote access administration. TACACS+:

- Provides three protocols (one each for authentication, authorization, and accounting). This allows each service to be provided by a different server.
- Uses TCP port 49.
- Encrypts the entire packet contents, not just authentication packets.
- Supports more protocol suites than RADIUS.

RADIUS is used by Microsoft servers for centralized remote access administration. RADIUS:

- Combines authentication and authorization using policies to grant access.
- Allows the separation of accounting to different servers. However, authentication and authorization remain combined on a single server.
- Uses UDP ports 1812 and 1813.
- Uses a challenge/response method for authentication. RADIUS encrypts only the password using MD5.

q_radius_tacacs_tac_dif_secp8

Which of the following are differences between RADIUS and TACACS+?

Answers:

- RADIUS uses TCP; TACACS+ uses UDP.
- RADIUS encrypts the entire packet contents; TACACS+ only encrypts the password.

- ***RADIUS combines authentication and authorization into a single function; TACACS+ allows these services to be split between different servers.**
- RADIUS supports more protocols than TACACS+.

Explanation:

TACACS+ provides three protocols (one each for authentication, authorization, and accounting). This allows each service to be provided by a different server. In addition, TACACS+:

- Uses TCP
- Encrypts the entire packet contents
- Supports more protocol suites than RADIUS

q_radius_tacacs_tac_port_secp8

Which of the following ports are used with TACACS?

Answers:

- 22
- ***49**
- 50 and 51
- 1812 and 1813
- 3389

Explanation:

Terminal Access Controller Access Control System (TACACS) uses port 49 for TCP and UDP.

Port 22 is used by Secure Shell (SSH).

Protocol numbers 50 and 51 are used by IPsec.

Ports 1812 and 1813 are used by Remote Authentication Dial-In User Service (RADIUS).

Port 3389 is used by Remote Desktop Protocol (RDP).

4.9 Network Authentication

As you study this section, answer the following questions:

- In the challenge/response process, what information is exchanged over the network during logon?
- What is included in a digital certificate?
- What is PKI?
- Which tool can manage authentication credentials on Windows hosts?

The key terms for this section include:

Term	Definition

Authentication	Authentication is the process of validating user credentials that prove user identity.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	3.8 Given a scenario, implement authentication and authorization solutions.
	• Authentication
	  ○ EAP
	  ○ 802.1X
	  ○ Single sign-on (SSO)
	  ○ Security Assertions Markup Language (SAML)
	  ○ OAuth
	  ○ OpenID
	  ○ Kerberos

TestOut Security Pro	1.0 Identity Management and Authentication 1.2 Harden Authentication

4.9.1 Network Authentication Protocols (Lesson Video)

Transcript:

Network authentication is the process of validating user credentials that prove user identity. Authentication is typically the first step in connecting to a network. Following successful authentication, access controls can be implemented to allow or deny access to network resources. In this lesson, we'll discuss single sign-on authentication, Federation, and OAuth—all methods used for network authentication.

Single Sign-on (SSO) is an authentication method that enables users to use the same login credentials to access multiple applications. A user only needs to authenticate once and then receive authorizations on compatible application servers without entering credentials again. Kerberos is a single sign-on network authentication and authorization protocol used on many networks, as implemented by Microsoft's Active Directory (AD) service.
Kerberos was named after the three-headed guard dog of Hades (Cerberus) because it consists of three parts: the client, the network resource, and a key distribution center. The clients request services from application servers, which rely on an intermediary—a key distribution center (KDC)—to vouch for their identity. Two services comprise a KDC: the Authentication Service and the Ticket Granting Service. Kerberos can authenticate human users and application services. These are collectively referred to as principals.
Using authentication to a Windows domain as an example:
The principal sends the authentication service (AS) a request for a Ticket Granting Ticket, or TGT. This is composed by encrypting the date and time on the local computer with the user's password hash as the key. The password hash itself isn't transmitted over the network. Although we refer to passwords for simplicity, the system can use other authenticators, such as smart card login.
The AS checks that the user account is present, that it can decode the request by matching the user's password hash with the one in the Active Directory database, and that the request hasn't expired. If the request is valid, the AS responds with a TGT and a TGS session key. A TGT contains information about the client (name and IP address) plus a time stamp and validity period. This is encrypted using the KDC's secret key. A TGS session key communicates between the client and the Ticket Granting Service (TGS). This is encrypted using a hash of the user's password.

The TGT is an example of a logical token. All the TGT does is identify who you are and confirm that you've been authenticated—it doesn't provide you with access to any domain resources. Presuming the user entered the correct password, the client can decrypt the Ticket Granting Service (TGS) session key but not the TGT. This establishes that the client and KDC know the same shared secret and that the client can't interfere with the TGT.

To access resources within the domain, the principal requests a service ticket—a token that grants access to a target application server. This process of granting service tickets is handled by the TGS. The principal sends the TGS a copy of its TGT, the name of the application server it wishes to access, and an authenticator consisting of a time-stamped client ID encrypted using the TGS session key. This confirms that the request is genuine. It also checks that the ticket hasn't expired or been used before.

The TGS service responds with a service session key and a service ticket. The service session key is used between the client and the application server. This is encrypted with the TGS session key. A service ticket contains information about the principal, such as a time stamp, system IP address, SID, the SIDs of groups to which it belongs, and the service session key. This is encrypted using the application server's secret key.

The principal forwards the service ticket, which it can't decrypt, to the application server and adds another time-stamped authenticator, which is encrypted using the service session key. The application server decrypts the service ticket to obtain the service session key using its secret key, confirming that the principal has sent it an untampered message. It then decrypts the authenticator using the service session key.

Now let's look at Federation. Federation is the notion that a network needs to be accessible to more than just a well-defined group of employees. A company might need to open parts of its network to partners, suppliers, and customers. The company can manage its employee accounts easily enough. Managing accounts for each supplier or customer internally may be more difficult. Federation means the company trusts accounts created and managed by a different network. As another example, in the consumer world, a user might want to use both Google Workspace and Twitter. Suppose Google and Twitter establish a federated network for authentication and authorization. In that case, users can log on to Twitter using their Google credentials or vice versa.

These interoperable federation protocols use claims-based identity. While the technical implementation and terminology differ, the overall model is like that of Kerberos SSO: The service provider, SP, and Identity Provider, or IdP, have a pre-established trust relationship. The user starts a session with the service provider.

The service provider redirects the principal to an IdP to authenticate. The user authenticates to the IdP. The IdP issues a claims token to the user. The user then presents the claims token to the SP. The SP validates that the IdP has signed the claim because of its trust relationship with the IdP. The service provider can now connect the authenticated principal to its accounts database to determine its permissions and other attributes.

Lastly, let's talk about Open Authorization. Open Authorization, often called OAuth, allows an application to authorize access without exposing the user's password. In a simplified sense, it's like a valet key for the web. It will enable a user's account information to be used by third-party services, such as Facebook or Twitter, without exposing their password. This means the user's sensitive details are handled only by the service they trust (i.e., their 'identity provider') and not the third-party applications requesting access.

OAuth is designed to facilitate the sharing of information or resources within a user profile between sites. The user creates a password-protected account at an identity provider. The user can link that identity to an OAuth consumer site without giving the password to the consumer site. Users can grant an OAuth client authorization to access some of their accounts. A client in this context is an app or consumer site. OAuth operates through a series of handshakes between three key players: the client—the application wanting access; the server—the application that holds the user's data; and the resource owner—the user.

Here's a simplified step-by-step of the process: The client requests authorization from the resource owner to access their data, which is held on the server. If the resource owner gives permission, the client receives an authorization grant. The client exchanges the authorization grant for an access token by requesting the server.

The server validates the authorization grant, and if it's correct, it issues an access token to the client. The client uses this access token to request the resource owner's data from the server. The server serves the client the requested data if the access token is valid. It's important to note that at no point does the client learn the resource owner's credentials, maintaining security and privacy of sensitive user information.

That's it for this lesson. In this lesson, we discussed network authentication. We talked about single sign-on authentication and authorization, federation, and open authorization. We also reviewed the step-by-step process for each of the network authentication methods.

4.9.2 Network Authentication Facts

This lesson covers the following topics:

- Network authentication overview
- Single sign-on authentication
- Single sign-on authorization
- Federation
- Open authorization

Network Authentication Overview

Authentication is the process of validating user credentials that prove user identity. Authentication is typically the first step in connecting to a network. Following successful authentication, access controls can be implemented to allow or deny access to network resources.

A simple form of authentication sends a username and password to an authentication server. If the password is sent in cleartext, the authentication credentials can be intercepted and used to impersonate an authorized user. One method of protecting login credentials is using a challenge/response mechanism (also called a three-way handshake). By doing so, both the authentication server and the authenticator are configured with a common shared secret. This shared secret is usually a password associated with a user account. The process is:

1. The authentication server sends a challenge string to the authenticator.
2. The authenticator uses the shared secret to hash the challenge string and returns the user account name and the hashed value to the authentication server.
3. The authentication server also uses its shared secret value to hash the challenge string. If the two hashed values match, the authentication server assumes the authenticator knows the shared secret.

Single Sign-on Authentication

A single sign-on (SSO) system allows users to authenticate once and then receive authorizations on compatible application servers without entering credentials again.

Kerberos is a single sign-on network authentication and authorization protocol used on many networks, as implemented by Microsoft's Active Directory (AD) service. Kerberos was named after the three-headed guard dog of Hades (Cerberus) because it consists of three parts. Clients request services from application servers, which rely on an intermediary—a key distribution center (KDC) —to vouch for their identity. Two services comprise a KDC: the Authentication Service and the Ticket Granting Service.

Kerberos can authenticate human users and application services. These are collectively referred to as principals. Using authentication to a Windows domain as an example, the first step in Kerberos SSO is to authenticate with a KDC server implemented as a domain controller.

- The principal sends the authentication service (AS) a request for a Ticket Granting Ticket (TGT). This is composed by encrypting the date and time on the local computer with the user's password hash as the key. The password hash itself is not transmitted over the network. Although we refer to passwords for simplicity, the system can use other authenticators, such as smart card login.
- The AS checks that the user account is present, that it can decode the request by matching the user's password hash with the one in the Active Directory database, and that the request has not expired. If the request is valid, the AS responds with the following data:
    - Ticket Granting Ticket (TGT) — contains information about the client (name and IP address) plus a time stamp and validity period. This is encrypted using the KDC's secret key.
    - TGS session key — communicates between the client and the Ticket Granting Service (TGS). This is encrypted using a hash of the user's password.

The TGT is an example of a logical token. All the TGT does is identify who you are and confirm that you have been authenticated—it does not provide you with access to any domain resources.

Single Sign-on Authorization

Presuming the user entered the correct password, the client can decrypt the Ticket Granting Service (TGS) session key but not the TGT. This establishes that the client and KDC know the same shared secret and that the client cannot interfere with the TGT.

- The principal requests a service ticket to access resources within the domain (a token that grants access to a target application server). This process of granting service tickets is handled by the TGS.
- The principal sends the TGS a copy of its TGT, the name of the application server it wishes to access, and an authenticator consisting of a time-stamped client ID encrypted using the TGS session key. The TGS should be able to decrypt both messages using the KDC's secret key for the first and the TGS session key for the second. This confirms that the request is genuine. It also checks that the ticket has not expired and has not been used before (replay attack).
- The TGS service responds with the following:
    - A Service session key — is used between the client and the application server. This is encrypted with the TGS session key.
    - A Service ticket — contains information about the principal, such as a time stamp, system IP address, Security Identifier (SID), the SIDs of groups to which it belongs, and the service session key. This is encrypted using the application server's secret key.
- The principal forwards the service ticket, which it cannot decrypt, to the application server and adds another time-stamped authenticator, which is encrypted using the service session key.
- The application server decrypts the service ticket to obtain the service session key using its secret key, confirming that the principal has sent it an untampered message. It then decrypts the authenticator using the service session key.
- Optionally, the application server responds to the principal with the time stamp used in the authenticator, which is encrypted using the service session key. The principal decrypts the time stamp, verifies that it matches the value already sent, and concludes that the application server is trustworthy. This means the server is authenticated to the principal (mutual authentication). This prevents an on-path attack, where a malicious user could intercept communications between the principal and server.
- The server now responds to access requests (assuming they conform to the server's access control list).

Federation

Federation is the notion that a network needs to be accessible to more than just a well-defined group of employees. A company might need to open parts of its network to partners, suppliers, and customers. The company can manage its employee accounts easily enough. Managing accounts for each supplier or customer internally may be more difficult. Federation means the company trusts accounts created and managed by a different network. As another example, in the consumer world, a user might want to use both Google Workspace and Twitter. Suppose Google and Twitter establish a federated network for authentication and authorization. In that case, users can log on to Twitter using their Google credentials or vice versa.

An on-premises network can use technologies such as LDAP and Kerberos, very often implemented as a Windows Active Directory network, because the administration of accounts and devices can be centralized. When implementing Federation, authentication and authorization design comes with more constraints and additional requirements to ensure interoperability between different platforms. Web applications might not support Kerberos, while third-party networks might not support direct Federation with Active Directory/LDAP. The design for these cloud networks likely requires using other standard protocols or frameworks for interoperability between web applications.

These interoperable federation protocols use claims-based identity. While the technical implementation and terminology are different, the overall model is similar to that of Kerberos SSO:

- The principal attempts to access a service provider (SP). The service provider redirects the principal to an identity provider (IdP) to authenticate.

- The principal authenticates with the identity provider and obtains a claim in the form of a token or document signed by the IdP.
- The principal presents the claim to the service provider. The SP can validate that the IdP has signed the claim because of its trust relationship with the IdP.
- The service provider can now connect the authenticated principal to its accounts database to determine its permissions and other attributes. It may be able to query attributes of the user account profile held by the IdP if the principal has authorized this type of access.

A federated network or cloud needs specific protocols and technologies to implement user identity assertions and transmit claims between the principal, the relying party, and the identity provider. Security Assertion Markup Language (SAML) is one such solution. SAML assertions (claims) are written in eXtensible Markup Language (XML). Communications are established using HTTP/HTTPS and the Simple Object Access Protocol (SOAP). The secure tokens are signed using the XML signature specification. Using a digital signature allows the relying party to trust the identity provider.

Open Authorization

Many public clouds use application programming interfaces (APIs) based on Representational State Transfer (REST) rather than SOAP. These are called RESTful APIs. Where SOAP is a tightly specified protocol, REST is a looser architectural framework. This allows the service provider more choice over implementation elements. Compared to SOAP and SAML, there is better support for mobile apps.

Authentication and authorization for a RESTful API are often implemented using the Open Authorization (OAuth) protocol. OAuth is designed to facilitate the sharing of information (resources) within a user profile between sites. The user creates a password-protected account at an identity provider (IdP). The user can link that identity to an OAuth consumer site without giving the password to the consumer site. A user (resource owner) can grant an OAuth client authorization to access some part of their account. A client in this context is an app or consumer site.

The user account is hosted by one or more resource servers. A resource server is called an API server because it hosts the functions that allow OAuth clients (consumer sites and mobile apps) to access user attributes. An authorization server processes authorization requests. A single authorization server can manage multiple resource servers; equally, the resource and authorization server could be the same server instance.

The client app or service must be registered with the authorization server. As part of this process, the client registers a redirect URL, which is the endpoint that will process authorization tokens. Registration also provides the client with an ID and a secret. The ID can be publicly exposed, but the client and the authorization server must keep the secret confidential. When the client application requests authorization, the user approves the authorization server to grant the request using an appropriate method. OAuth supports several grant types—or flows—for use in different contexts, such as server to server or mobile app to server. Depending on the flow type, the client will end up with an access token validated by the authorization server. The client presents the access token to the resource server, which then accepts the request for the resource if the token is valid.

4.9.3 LDAP Authentication (Lesson Video)

Transcript:

In this lesson, we'll discuss the Lightweight Directory Access Protocol authentication, or LDAP.

LDAP is an open-source protocol used to communicate with network directories. It is a lightweight, fast protocol that runs over TCP/IP. That makes it ideal for internet-based access requests. Because LDAP is an open protocol, most applications can access a desired server regardless of the directory service being used to manage authentication. LDAP traffic is in cleartext by default, but LDAP can run using SSL or some type of transport-level security.
So, how exactly does LDAP work? Let's say we have a client who is trying to access network resources. The network's server has been configured with Active Directory.

Active Directory includes a database with details about authorized network user's information including usernames, passwords, and groups. It also stores information about various network resources including the level of access each user has to that resource.
To be compliant with the LDAP protocol, data must be stored using a standard method. For example, each database entry must be an object with specific attributes. This structure helps to ensure that a directory can be easily accessed and the desired information can be found.
To obtain access to the network, Active Directory and the client must speak the same language. This is where LDAP comes in. The client sends a request with required credentials to Active Directory. When Active Directory receives the request and the credentials, it compares them the information to the database. If the information matches, the request is authorized. The client is then connected to the appropriate network resources.
You must set up LDAP to authenticate credentials against the information stored in directory services. A bind method sets the authentication state for an LDAP session. You can use one of two options to perform the authentication.
Simple Authentication uses name and password authentication, unauthenticated authentication, or anonymous authentication. Most of the time, a name and a password create the bind request that's sent to the server.
Simple Authentication and Security Layer authentication, or SASL, uses a different authentication service such as Kerberos to bind to the LDAP server. The LDAP server then uses this authentication service. This method provides additional security because the authentication method is separated from the application protocols.
As you can see, while directory services provide security specific to your network, LDAP provides a method for communication with remote clients. LDAP and directory services work hand in hand to keep your network secure.
That's it for this lesson. In this video, we discussed the Lightweight Directory Access Protocol. We discussed what it is and how it works. We also covered the two authentication methods, simple authentication and SASL authentication.

4.9.4 LDAP Authentication Facts

This lesson covers the following topics:

- Directory services
- Securing directory services

Directory Services

A directory service stores information about users, computers, security groups/roles, and services. Each object in the directory has several attributes. The directory schema describes the types of attributes, what information they contain, and whether they are required or optional. For products from different vendors to be interoperable, most directory services are based on the Lightweight Directory Access Protocol (LDAP), which was developed from a standard called X.500.

Within an X.500-like directory, a distinguished name (DN):

- Is a collection of attributes that define a unique identifier for any given resource.
- Is made up of attribute-value pairs, separated by commas. The most specific attribute is listed first, and successive attributes become progressively broader. This most specific attribute is the relative distinguished name, as it uniquely identifies the object within the context of successive (parent) attribute values.

Some of the attributes commonly used include common name (CN), organizational unit (OU), organization (O), country (C), and domain component (DC).

CN=WIDGETWEB, OU=Marketing, O=Widget, C=UK, DC=widget, DC=foo

Securing Directory Services

A network directory lists the subjects (principally users, computers, and services), objects (such as directories and files) available on the network, and subjects' permissions over objects. A directory facilitates authentication and authorization and must be maintained as a highly secure service. Most directory services are based on the Lightweight Directory Access

Protocol (LDAP), running over port 389. The basic protocol provides no security, and all transmissions are plaintext, making it vulnerable to sniffing and on-path attacks. Authentication (referred to as binding to the server) can be implemented in the following ways:

- No Authentication — means anonymous access is granted to the directory.
- Simple Bind — means the client must supply its distinguished name (DN) and password, passed as plaintext.
- Simple Authentication and Security Layer (SASL) — means the client and server negotiate using a supported authentication mechanism, such as Kerberos. The STARTTLS command can require encryption (sealing) and message integrity (signing). This is the preferred mechanism for Microsoft's Active Directory (AD) implementation of LDAP.
- LDAP Secure (LDAPS) — means the server is installed with a digital certificate, which it uses to set up a secure tunnel for the user credential exchange. LDAPS uses port 636.

If secure access is required, anonymous and simple authentication methods should be disabled on the server.

Generally, two access levels must be granted on the directory: read-only access (query) and read/write access (update). This is implemented using an access control policy, but the precise mechanism is vendor-specific and not specified by the LDAP standards documentation.

Unless hosting a public service, the LDAP directory server should only be accessible from the private network. This means that the LDAP port should be blocked by a firewall from accessing the public interface. If there is integration with other services over the internet, ideally, only authorized IPs should be permitted.

4.9.5 Practice Questions (Section Quiz)

q_netauth_federation_secp8

A manufacturing company recently bought out another similar company. They need to link each company's directory systems together to access their resources without merging the two.

How can they link the two directory systems together?

Answers:

- Site-to-site VPN
- Migration
- ***Federation**
- Location-based restrictions

Explanation:

Federation directories allow two different subsets of accounts to work together for permissions and access.

A site-to-site virtual private network (VPN) is similar to the VPN that employees use. However, they connect two or more location-separated corporate resources.

Over time, merging companies will migrate employee accounts from one directory to another. However, they do not need to work with each other seamlessly.

Location-based restrictions are not as efficient and do not provide for linking directories together.

q_netauth_kerberos_secp8

When using Kerberos authentication, which of the following terms is used to describe the token that verifies the user's identity to the target system?

Answers:

- Coupon
- Voucher
- ***Ticket**
- Hashkey

Explanation:

The tokens used in Kerberos authentication are known as *tickets*. Tickets perform a number of functions, including notifying the network service of the user who has been granted access and authenticating the identity of that person when he or she attempts to use the network service.

The terms *coupon* and *voucher* are not associated with Kerberos or any other commonly implemented network authentication system.

The term *hashkey* is sometimes used to describe a value that has been derived from some piece of data when that value is then used to access a service. This term is not associated with Kerberos.

q_netauth_mutual_secp8

What is mutual authentication?

Answers:

- ***A process by which each party in an online communication verifies the identity of the other party.**
- The use of two or more authentication factors.
- Deploying CHAP and EAP on remote access connections.
- Using a certificate authority (CA) to issue certificates.

Explanation:

Mutual authentication is the process by which each party in an online communication verifies the identity of the other party. Mutual authentication is most common in VPN links, SSL connections, and e-commerce transactions. In each of these situations, both parties in the communication want to ensure that they know with whom they are interacting.

The use of two or more authentication factors is called two-factor authentication. Challenge Handshake Authentication Protocol (CHAP) and Extensible Authentication Protocol (EAP) are authentication protocols.

Communicating hosts might use certificates issued by a trusted CA in performing mutual authentication. However, using the CA is not, in itself, a definition of mutual authentication.

q_netauth_oauth_01_secp8

In a company, different departments actively access various cloud-based applications and services to perform their tasks efficiently. The company's security team has concerns about the growing complexity and risks of managing user credentials across multiple platforms.

To address this concern proactively, the team implements a modern authentication solution that actively provides single sign-on (SSO) capabilities, ensuring enhanced user convenience and security.

In this scenario, which technology should the organization proactively employ for federation and enabling SSO capabilities effectively across the diverse range of cloud-based applications?

Answers:

- ***Open Authorization (OAuth)**
- Role-based access control (RBAC)
- Lightweight Directory Access Protocol (LDAP)
- Public key infrastructure (PKI)

Explanation:

In this scenario, the organization uses Open Authorization (OAuth) for federation, allowing secure authorization and delegation of user access to third-party applications without exposing user credentials.

While role-based access control (RBAC) is a valuable access control model, it is not directly related to federation and SSO capabilities across multiple cloud-based applications.

Although used to access and manage directory information, Lightweight Directory Access Protocol (LDAP) is not the technology used for federation and SSO in the given scenario.

Public key infrastructure (PKI) is a cryptographic system that uses public and private keys for secure communication and authentication. While it plays a role in security, there are other technologies for federation and SSO in this context.

q_netauth_oauth_02_secp8

In a large healthcare organization, multiple departments handle sensitive patient data. Each department requires access to different applications and systems to carry out its tasks efficiently.

However, granting broad access rights through long-lived authentication tokens poses security risks.

What solution should the IT department implement while adhering to the principle of least privilege and securing sensitive patient data?

Answers:

- ***Open Authorization (OAuth)**
- JSON web token (JWT)
- Multi-factor authentication (MFA)
- Kerberos

Explanation:

OAuth (Open Authorization) is a widely used authentication framework that enables secure authorization between different services and aligns with the principle of least privilege.

JWT (JSON web token) is a compact, URL-safe means of representing claims between two parties. Unlike OAuth, which focuses on authorization delegation, JWT is more about transmitting claims.

MFA (Multi-factor authentication) enhances security by requiring users to provide multiple forms of verification before accessing systems. While MFA is valuable, it doesn't inherently address the issue of granting time-limited access to different departments.

Kerberos is a network authentication protocol that relies on tickets for granting access. While it offers secure authentication, it doesn't directly address the issue of temporal or ephemeral access rights.

q_netauth_oauth_03_secp8

A company wants to set up single sign-on (SSO) without passing credentials through to each piece of software and cloud service.

Which protocol would meet this requirement?

Answers:

- Kerberos
- FIDO
- VPN
- ***OAuth**

Explanation:

The Open Authorization (OAuth) protocol is a system that facilitates sharing of information (resources) within a user profile between sites. The user can link that identity to an OAuth consumer site without giving the password to the consumer site.

The preferred system for network authentication in a Windows environment is Kerberos which replaces the legacy system NT LAN Manager (NTLM) authentication.

Passwordless authentication means that the whole system no longer processes knowledge-based factors. The Fast IDentity Online 2 (FIDO2) with WebAuthn specifications provides a framework for passwordless authentication.

A virtual private network (VPN) is how an individual can access corporate resources outside the corporate infrastructure.

q_netauth_oauth_04_secp8

You are a security analyst at a large corporation. Your company has recently implemented an Open Authorization (OAuth) system to allow third-party applications to access company resources.

One day, you notice an unusual amount of data being transferred from your company's servers to an unknown third-party application.

What should be your first course of action?

Answers:

- Ignore the situation as the OAuth system is designed to allow third-party applications to access company resources.
- Immediately block all data transfers to the third-party application without further investigation.
- ***Investigate the third-party application to understand why it is accessing a large amount of data.**
- Report the situation to the company's legal department without conducting any further investigation.

Explanation:

Investigating the third-party application is the correct answer. This action will help you understand why the OAuth system is accessing a large amount of data. This could reveal whether the data transfer is legitimate or if there is a potential security breach.

While OAuth is designed to allow third-party applications to access company resources, an unusual amount of data transfer could indicate a potential security breach or misuse of access privileges.

While it might be necessary to block data transfers eventually, doing so immediately without further investigation could disrupt legitimate business operations and does not address the root cause of the issue.

While it might be necessary to involve the legal department eventually, doing so without conducting any further investigation does not address the immediate potential security risk.

q_netauth_saml_01_secp8

In a multinational corporation, employees across various departments regularly access many cloud-based applications to fulfill their tasks efficiently. The company's security team is grappling with managing user credentials securely and efficiently across these diverse platforms.

They are actively looking to improve user authentication and streamline access to these applications while ensuring robust security measures are in place.

In this scenario, what technology should the company implement to enable single sign-on (SSO) capabilities and ensure secure authentication across its diverse cloud-based applications?

Answers:

- ***Security Assertion Markup Language (SAML)**
- Lightweight Directory Access Protocol (LDAP)
- Remote Authentication Dial-In User Service (RADIUS)
- Virtual private network (VPN)

Explanation:

Security Assertion Markup Language (SAML) enables secure SSO across various applications by exchanging authentication and authorization data between parties through an extensible markup language (XML)-based protocol.

Lightweight Directory Access Protocol (LDAP) queries and modifies directory services but does not work with SSO or cloud-based application authentication.

Remote Authentication Dial-In User Service (RADIUS) is primarily utilized for authenticating users in network access scenarios and does not cater to SSO across cloud-based applications.

A virtual private network (VPN) establishes secure connections between remote networks, but there are no links to enabling SSO or cloud-based application authentication.

q_netauth_saml_02_secp8

A real estate investment firm wants to implement single sign-on (SSO) for its dozens of services and software. The firm found a vendor to implement that request using the eXtensible Markup Language (XML) standard.

What solution does this vendor use for SSO?

Answers:

- ***SAML**
- VPN
- LDAP
- LSASS

Explanation:

Security Assertion Markup Language (SAML) allows for federating a network or cloud system. SAML assertions and claims between the principal, the relying party, and the identity provider use eXtensible Markup Language as their structure.

A virtual private network (VPN) allows individuals to access corporate resources outside the corporate infrastructure. Employees cannot access a file share without a VPN if they are not at their office.

LDAP is a protocol that makes it possible for applications to query user information rapidly. It is not an XML solution.

LSASS is a process in Microsoft Windows operating systems that is responsible for enforcing the security policy on a system.

q_netauth_shared_secret_secp8

In a three-way handshake during network authentication, what is the term for the common value that both the authentication server and the authenticator are configured with?

Answers:

- ***Shared secret**
- Public key
- Private key
- Digital signature

Explanation:

Shared secret is the correct answer. In a three-way handshake, both the authentication server and the authenticator are configured with a common value known as a shared secret. This is usually a password associated with a user account.

In public key cryptography, the public key is used to encrypt data and it's publicly available. It's not the common value used in a three-way handshake.

In public key cryptography, the private key is kept secret by the owner and is used to decrypt data. It's not the common value used in a three-way handshake.

A digital signature is a mathematical scheme for verifying the authenticity of digital messages or documents. It's not the common value used in a three-way handshake.

q_netauth_tgs_secp8

In a Kerberos authentication system, how does the Ticket Granting Service (TGS) contribute to the single sign-on (SSO) process?

Answers:

- The TGS generates the initial Ticket Granting Ticket (TGT) for the client.
- The TGS validates the client's password and username.
- ***The TGS issues service tickets to clients for accessing specific services.**
- The TGS encrypts all data transferred between the client and the application server.

Explanation:

The TGS issues service tickets to clients after they have been authenticated. These service tickets allow clients to access specific services without having to re-authenticate.

The initial Ticket Granting Ticket (TGT) is generated by the authentication service (AS), not the TGS.

The validation of the client's password and username is also done by the authentication service (AS), not the TGS.

While the TGS does play a role in the encryption process by providing session keys, it does not encrypt all data transferred between the client and the application server. The encryption of data transferred between the client and the server is done using session keys, which are provided by the TGS but the actual encryption is done at the client and server ends.

q_ldap_access_levels_secp8

Which of the following are the access levels that are generally granted on the directory in LDAP? (Select two.)

Answers:

- ***Read-only access**
- ***Read/write access**
- Full control access
- Execute access
- Delete access

Explanation:

The following are the access levels generally granted on the directory in LDAP:

- Read-only access (query) - This level of access allows users to view and query the data in the directory but not modify it. It's essential for users who need to retrieve information but should not change it.
- Read/write access (update) - This level of access allows users to both view and modify the data in the directory. It's necessary for users who need to update or change the information in the directory.

While full control access is used in some systems to denote a user having all possible permissions, it's not a standard access level in LDAP. In LDAP, permissions are typically more granular and are defined by the specific access (read, write, etc.) a user has to specific objects or attributes.

Execute access is typically used in file system permissions to denote the ability to run a file as a program. In the context of LDAP, which is a directory service, this concept doesn't apply.

While the ability to delete objects is a type of permission that can be granted in LDAP, it's not one of the two primary access levels that must be granted. The primary levels are read-only and read/write.

q_ldap_dn_secp8

What BEST describes a distinguished name (DN) in the context of a directory service?

Answers:

- ***A unique identifier for any given resource in a directory, made up of attribute-value pairs.**
- A security protocol used to protect directory services.
- A type of directory service developed from the X.500 standard.
- A method of authentication in the Lightweight Directory Access Protocol (LDAP).

Explanation:

A distinguished name (DN) is a unique identifier for any given resource in a directory, made up of attribute-value pairs. The DN provides a way to locate a specific object in the directory.

A DN is not a security protocol. It is a unique identifier used in directory services.

While DN is used in directory services, it is not a type of directory service itself. X.500 is a standard from which LDAP, a type of directory service, was developed.

A DN is not a method of authentication. It is a unique identifier used in directory services. Authentication methods in LDAP include No Authentication, Simple Bind, Simple Authentication and Security Layer (SASL), and LDAP Secure (LDAPS).

q_ldap_firewall_secp8

In the context of LDAP security, why would an organization choose to block the LDAP port with a firewall?

Answers:

- ***To prevent unauthorized access to the directory from the public network.**
- To disable the LDAP service entirely.
- To force the use of LDAPS instead of LDAP.
- To prevent the directory from storing any new data.
- To disable all network communication.

Explanation:

The correct answer is to prevent unauthorized access to the directory from the public network. Blocking the LDAP port with a firewall can help prevent unauthorized access from the public network. This is a common security measure to protect sensitive directory information.

To disable the LDAP service entirely is incorrect. Blocking the LDAP port with a firewall does not disable the LDAP service. It simply restricts access to the service from certain networks.

While LDAPS (LDAP Secure) does use a different port than standard LDAP, blocking the LDAP port does not force the use of LDAPS. The use of LDAPS must be configured separately.

Blocking the LDAP port with a firewall does not prevent the directory from storing new data. It only restricts network access to the LDAP service, not all network communication. Other services and protocols can still communicate over the network.

q_ldap_ou_secp8

Which of the following is a commonly used attribute in Lightweight Directory Access Protocol (LDAP)?

Answers:

- ***Organizational Unit (OU)**
- Internet Protocol (IP)
- Firewall
- Public network

Explanation:

Organizational Unit (OU) is the correct answer. In LDAP, an Organizational Unit (OU) is a type of attribute used to organize objects within a directory. It can represent a logical or physical entity in the organization, such as a department or a location.

While IP addresses are crucial for network communication, they are not considered an attribute in LDAP. LDAP attributes are characteristics of directory objects, such as users or devices, not network protocols.

A firewall is a network security device that monitors and filters incoming and outgoing network traffic based on an organization's previously established security policies. It is not an attribute in LDAP.

A public network refers to a network that is accessible to the general public. It is not an attribute in LDAP. LDAP attributes are used to define characteristics of objects within the directory, not types of networks.

q_ldap_sasl_secp8

An educational institution's systems administrator is responsible for securing the LDAP directory service for the organization's computing resources.

Which authentication method should the systems administrator implement to ensure secure access?

Answers:

- No authentication
- Simple Bind
- ***Simple Authentication and Security Layer (SASL)**
- LDAP Secure (LDAPS)

Explanation:

SASL allows the client and server to negotiate a supported authentication mechanism and provides the option to use the command STARTTLS for encryption and message integrity. This feature is a secure way to access the Lightweight Directory Access Protocol (LDAP) directory.

Enabling anonymous access (no authentication) to the directory is not secure, exposing it to potential misuse and unauthorized access.

In a simple bind, the client either binds anonymously, that is, with an empty bind DN, or by providing a DN and a password. This is not a secure method for LDAP.

LDAP Secure (LDAPS) uses TLS/SSL as a transmission protocol. However, this is not as secure as using SASL.

q_ldap_secure_secp8

The IT administrator for a large university uses an LDAP directory service to manage user access to various computing resources.

To ensure the directory's security, which of the following measures should the administrator implement?

Answers:

- Allow anonymous access to the directory for easy user onboarding.
- Use the basic LDAP protocol without any additional security mechanisms.
- Implement Simple Bind with plaintext transmission of distinguished name and password.
- ***Set up LDAP Secure (LDAPS) with a digital certificate on port 636 for secure user credential exchange.**

Explanation:

Setting up LDAP Secure (LDAPS) with a digital certificate on port 636 for secure user credential exchange encrypts data and ensures the protection of user credentials during transmission by providing a secure tunnel.

Exposing sensitive data and creating potential security risks results from anonymous access to the directory, making it not recommended for secure environments.

Using the basic LDAP protocol without any additional security mechanisms does not provide the extra security needed in this scenario.

Implementing Simple Bind with plaintext transmission of distinguished name and password does not provide the level of security needed in this scenario.

5.0 Network Architecture

5.1 Enterprise Network Architecture

As you study this section, answer the following questions:

- What is network architecture?
- What is Internet Protocol?
- What needs should be considered when setting up security zones?
- What security protocols help reduce the attack surface?

The key terms for this section include:

Term	Definition
Network architecture	The selection and placement of media, devices, protocols/services, and data assets.
Network infrastructure	The media, appliances, and addressing/forwarding protocols that support basic connectivity.
Internet Protocol (IP)	Provides the addressing mechanism for logical networks and subnets.
Attack surface	All the points at which a threat actor could gain access to hosts and services.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	3.1 Compare and contrast security implications of different architecture models.
	• Architecture and infrastructure concepts
	  ○ Cloud
	  ▪ Responsibility matrix
	  ▪ Hybrid considerations
	  ▪ Third-party vendors
	  ○ Infrastructure as code (IaC)
	  ○ Serverless
	  ○ Microservices
	  ○ Network infrastructure
	  ▪ Physical isolation
	  ▪ Air-gapped
	  ▪ Logical segmentation
	  ▪ Software-defined networking (SDN)
	  ▪ On-premises

	▪ Centralized/decentralized • Considerations ○ Availability ○ Resilience ○ Cost ○ Responsiveness ○ Scalability ○ Ease of deployment ○ Risk transference ○ Ease of recovery ○ Patch availability ○ Inability to patch ○ Power ○ Compute 3.2 Given a scenario, apply security principles to secure enterprise infrastructure. • Infrastructure considerations ○ Device placement ○ Security zones ○ Attack surface ○ Connectivity ○ Failure modes ▪ Fail-open ▪ Fail-closed ○ Device attribute ▪ Active vs. passive ▪ Inline vs. tap/monitor ○ Load balancer • Selection of effective controls

5.1.1 Enterprise Network Architecture (Lesson Video)

Transcript:

In this lesson, we'll explore the fundamentals of network design, from architecture decisions to device placement. Let's get started!

Network architecture is like the blueprint of a network. It involves choosing and arranging components such as media, devices, protocols or services, and data assets. Within this framework, we have network infrastructure, which consists of the physical elements like cables, appliances, and addressing and forwarding protocols that enable basic connectivity. Network applications are the services running on this infrastructure to support business activities like processing invoices or sending emails. Data assets are the information generated as a result of these activities.
The goal of secure network infrastructure and application architecture is to support secure business workflows, where security means ensuring confidentiality, integrity, and availability.
Let's take the example of provisioning email services to understand the architecture decisions involved.
Access control ensures that client devices access the network through a secure channel with authentication and authorization. Unauthorized users and devices must be denied access.
The email mailbox server stores data assets securely and must be accessible only to authorized clients. It should also be fault-tolerant and highly available.

The mail transfer server, responsible for connecting with untrusted internet hosts, requires careful control to manage communications between the untrusted network and trusted LAN. Any data or software leaving or entering the network must be subject to policy-based controls.
This type of business flow will involve systems with different security requirements. Placing the client, the mailbox, and the mail transfer server all within the same segment will introduce many vulnerabilities. Understanding and controlling how data flows between these network segments is a key part of secure and effective network architecture design.
Now, let's dive into network infrastructure using the OSI model, a framework for understanding network layers.
At the physical layer (Layer 1), we find various transmission media, from cables to wireless signals, which are the links connecting network nodes. Now, nodes come in two types:
host nodes, which initiate data transfers (like servers or clients), and intermediary nodes that forward traffic around the network.
Networks can vary in scope, from local area networks (LANs) for a single site to wide area networks (WANs) that span metropolitan, country-wide, or global scopes.
Each network node must have a unique address, and this addressing function occurs at different layers and scopes. Various network appliances and protocols handle forwarding and addressing functions, such as switches, routers, transport protocols, application protocols, and DNS servers.
When designing a network, several factors come into play. Costs, including upfront investments and ongoing maintenance, need to be considered. Compute resources affect processing time for various workloads. Scalability and ease of deployment help manage changes in workload. Availability and resilience minimize downtime and recovery time in case of failures. Facility power usage capabilities and patch availability are also vital considerations, along with the option of risk transference through service-level agreements. These factors together shape the architecture of your network.
The selection of effective controls for network infrastructure is the process of choosing the type and placement of security appliances and software. The aim is to enforce segmentation, apply access controls, and monitor traffic for policy violations.
Guided by the principle of defense in depth, diverse controls are placed at different OSI layers to enhance security. You have three options: preventive controls at network borders, detective controls within the perimeter, and a combination of preventive, detective, and corrective controls on hosts. For instance, a firewall at the network border enforces access rules, while a sensor behind it relays traffic to an intrusion detection system for monitoring. Access control lists on internal routers enforce rules for traffic between zones, and load balancers can help mitigate denial-of-service attacks.
Endpoint protection software on hosts adds another layer of security, implementing host firewalls, anti-virus, intrusion detection, and data loss prevention.
Lastly, attributes determine the precise way in which a device can be placed within the network topology.
Controls can be active or passive. Active controls require configuration and data exchange, while passive controls operate silently.
Deployment can be inline, becoming part of the network path, or as taps or monitors, capturing traffic transparently. The router/firewall is an active control as client devices must be configured to use it for internet access. The TAP and mirror ports are passive controls. They're completely transparent to the server and client hosts.
That's it for this lesson. In this lesson, we discussed network architecture and how it's the blueprint of your network design. We then talked about how the network infrastructure fits within the architecture framework. We used the OSI model to help us understand how different components of the infrastructure fit together. Next, we reviewed the many architecture considerations that factor into your network design. We also discussed device placement and security controls. We finished this lesson by looking at device attributes that determine how a security control affects device placement.

5.1.2 Enterprise Network Architecture Facts

This lesson covers the following topics:

- Network architecture
- Network infrastructure and OSI model
- Architecture considerations
- Device placement and security controls
- Device attributes

Network Architecture

Network architecture means the selection and placement of media, devices, protocols/services, and data assets:

- Network infrastructure is the media, appliances, and addressing/forwarding protocols that support basic connectivity.
- Network applications are the services that run on the infrastructure to support business activities, such as processing invoices or sending email.
- Data assets are the information that is created, stored, and transferred as a result of business activity.

Secure network infrastructure and application architecture are put there to support secure business workflows. A workflow is a series of tasks that a business needs to perform, such as accepting customer orders from a web store. Remember that security means the attributes of confidentiality, integrity, and availability.

Analyzing the systems involved in provisioning email can illustrate the sorts of architecture decisions that need to be made:

- Access—the client device must access the network via a physical channel and obtain a logical address. The user must be authenticated and authorized to use the email application. The corollary is that unauthorized users and devices must be denied access.
- Email mailbox server—the mailbox stores data assets and must only be accessed by authorized clients, and conversely, must be fully available and fault tolerant to support the genuine user. The email service must run with a minimum number of dependencies over network infrastructure that is resilient to faults.
- Mail transfer server—this must connect with untrusted internet hosts, so communications between the untrusted network and trusted LAN must be carefully controlled. Any data or software leaving or entering the network must be subject to policy-based controls.

This type of business flow will involve systems with different security requirements. Placing the client, the mailbox, and the mail transfer server all within the same segment will introduce many vulnerabilities. Understanding and controlling how data flows between these network segments is a key part of secure and effective network architecture design.

Network Infrastructure and the OSI Model

It is helpful to use a layer model to analyze network infrastructure and services. The Open Systems Interconnection (OSI) model is a widely quoted example of how to define layers of network functions.

A network is comprised of nodes and links. At the physical (PHY) layer, or layer 1 in the OSI model, links are implemented as twisted-pair cables transmitting electrical signals, fiber optic cables carrying infrared light signals, or as wireless devices transmitting radio waves.

There are two types of nodes. A host node is one that initiates data transfers. Hosts are usually either servers or clients. An intermediary node forwards traffic around the network. This forwarding occurs at different layers and with different scopes. A network scope of a single site is referred to as a local area network (LAN). Networks that span metropolitan, country-wide, or global scopes are called wide area networks (WANs).

Each network node must be identifiable via a unique address. This addressing function also takes place at different layers with different scopes.
Forwarding and addressing functions are handled by the following network appliances and protocols:

- A switch forwards frames between nodes in a cabled network. The network adapter in each host is connected to a switch port via a cable. Switches work at layer 2 of the OSI model. Most LANs use networks based on the Ethernet standard. An Ethernet switch makes forwarding decisions based on the hardware or media access control (MAC) address of attached hosts. A MAC address is a 48-bit value written in

hexadecimal notation, such as00-15-5D-01-CA-4A. This addressing works within the local network segment only. This is referred to as a broadcast domain.

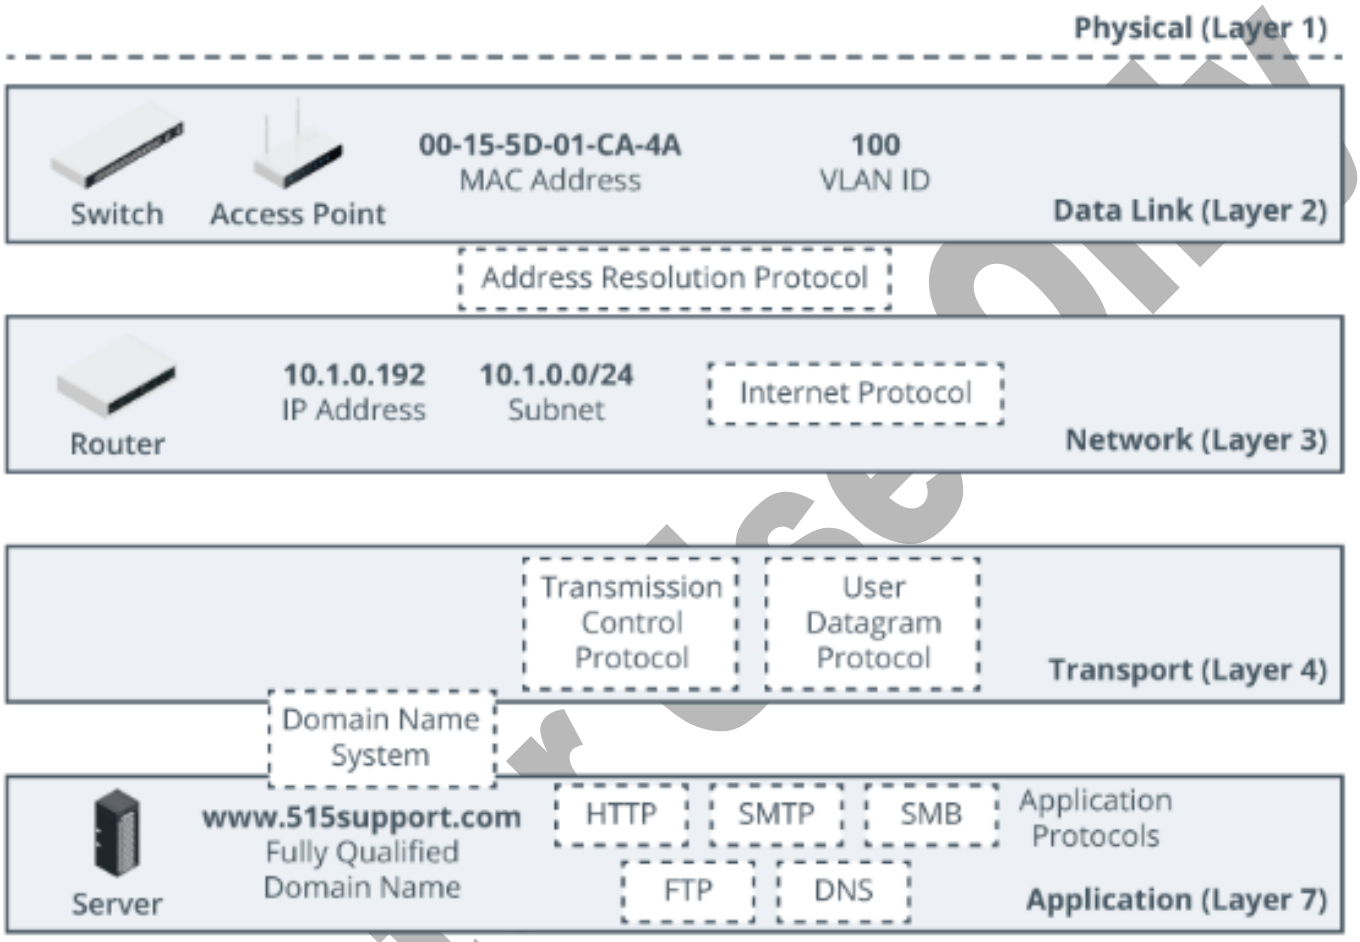

Appliances, protocols, and addressing functions within the OSI network layer reference model. (Images © 123RF.com.)

- Wireless access points provide a bridge between a cabled network and wireless hosts, or stations. Access points work at layer 2 of the OSI model. Wireless devices also use MAC addressing at layer 2.
- Routers send packets around an internetwork, making forwarding decisions based on Internet Protocol (IP) addresses. Routers work at layer 3 of the OSI model. Each local segment will normally have a router connected to it. The router acts as a default gateway for hosts on the segment to use to send packets to other segments.
- Transport protocols allow clients to exchange data with application servers. The Transmission Control Protocol (TCP) establishes reliable connections, while the User Datagram Protocol (UDP) allows unreliable, connectionless transfers. Each application protocol is identified by a TCP or UDP port. This functionality is defined at layer 4 of the OSI model.
- Application protocols support client/server functionality for user-level services, such as web browsing, email, and file transfer. Application protocols work at layer 7 of the OSI model.
- Domain Name System (DNS) servers host name records and perform name resolution to allow applications and users to address hosts and services using fully qualified domain names (FQDNs) rather than IP addresses. DNS also works at layer 7 of the OSI model, but is an infrastructure service, rather than a user-level service, like web browsing.

The OSI model has three upper layers. In practical terms, distinguishing the functions of layers 5, 6, and 7 isn't that helpful, so just think of applications working at layer 7.

Architecture Considerations

When evaluating the use of a particular architecture and selecting effective controls, consider a number of factors:

- Costs—architecture changes, the acquisition and upgrade of appliances, and software require an up-front capital outlay, which can depreciate and lose value. There are also ongoing maintenance and support liabilities. The value of the investment in security architecture and controls can be calculated based on how much they reduce losses from incidents.
- Compute and responsiveness—minimize processing time for workloads. A workload is the processing effort required to complete a task, such as a web server responding to a client request. Each network device requires sufficient CPU, system memory, storage, and network bandwidth resources to ensure an acceptable response time for a given workload. Higher compute resources incur greater costs.
- Scalability and ease of deployment —minimize costs when workloads increase or decrease. If workloads decrease it can be difficult to recover capital costs. If workloads increase it can be difficult to deploy new nodes or upgrade existing nodes to maintain responsiveness. A scalable system is one that can quickly or automatically add or remove compute resources without incurring excessive costs.
- Availability—minimizes downtime or maximizes uptime. Downtime represents the loss of opportunity to do work, damaging the business's reputation, revenue, and profitability. Downtime can be due to planned maintenance or unplanned failures and security incidents.
- Resilience and ease of recovery —reduce the time to recover from a failure. For example, a system that recovers from a failure without manual intervention is more resilient than one that requires an administrator to restart it.
- Power—is a feature that ensures the facility can meet the energy demands of its devices and workloads. Power usage through higher compute resources increases costs. Ensuring that the building infrastructure minimizes power failures improves availability.
- Patch availability—ensures that firmware and software code is protected against exploits for known vulnerabilities. Conversely, the network owner cannot manage this process when they rely on a third party to maintain infrastructure or when a device or software product is no longer supported by its vendor.
- Risk transference—is a contract that uses a third party to manage the network infrastructure. A service-level agreement (SLA) can be defined with penalties if metrics for responsiveness, scalability, availability, and resilience are not maintained.

On-premises networks tend to have high capital costs and low scalability. For example, consider the difficulty of increasing bandwidth from 1 Gbps to 10 Gbps operation across the entire site. This would likely require the installation of new cable throughout the building. Recovery procedures can be complex if the site premises is affected by a large-scale disaster. This means that availability and resilience can be lower than alternative solutions such as cloud networking.

Device Placement and Security Controls

The selection of effective controls for network infrastructure is the process of choosing the type and placement of security appliances and software. The aim is to enforce segmentation, apply access controls, and monitor traffic for policy violations. The selection of effective controls is governed by the principle of defense in depth. Defense in depth means that security-critical zones are protected by diverse preventive, detective, and corrective controls operating at each layer of the OSI model. Defense in depth is ensured through careful selection of device placement within the network topology. There are three options:

- Preventive controls—are often placed at the border of a network segment or zone. Preventive controls such as firewalls enforce security policies on traffic entering and exiting the segment, ensuring confidentiality and integrity. A load balancer control ensures high availability for access to the zone.
- Detective controls—might be placed within the perimeter to monitor traffic exchanged between hosts within the segment. This provides alerting of malicious traffic that has evaded perimeter controls.

- Preventive, detective, and corrective controls —might be installed on hosts as a layer of endpoint protection in addition to the network infrastructure controls.

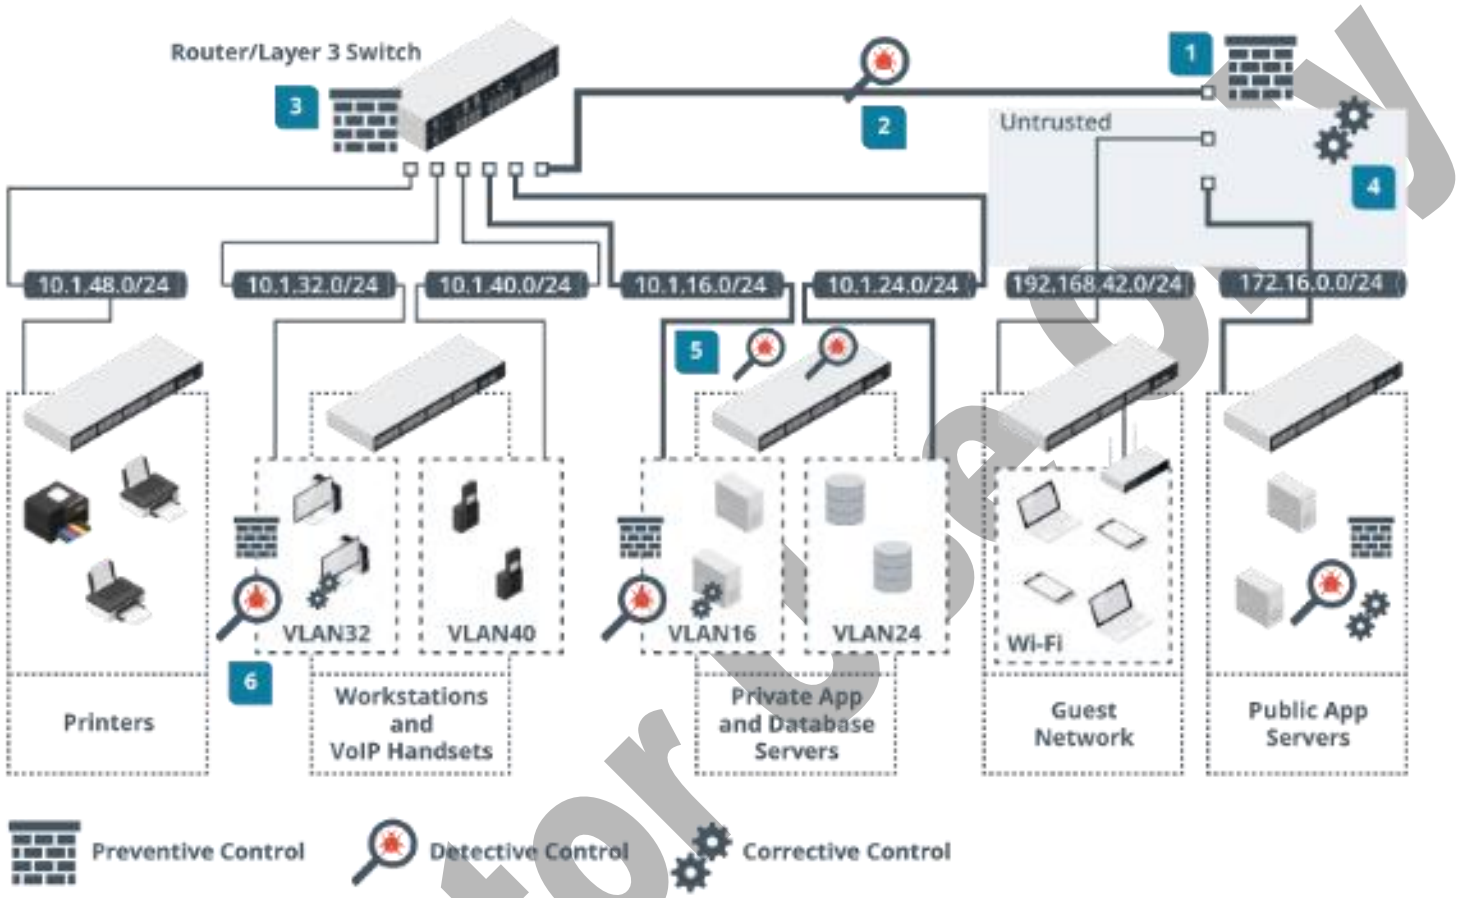

Placement of security controls to ensure diversity and defense in depth.

As an illustration, the diagram shows how different control types can be positioned within the network to ensure defense in depth:

- At the network border, a preventive control such as a firewall enforces access rules for ingress and egress traffic.
- A sensor placed inline behind the border firewall relays traffic to an intrusion detection system to implement detective control and identify malicious traffic that has evaded the firewall.
- Access control lists configured on internal routers enforce rules for traffic being forwarded between internal zones and hosts.
- Incoming traffic for public-facing servers can be mediated by a load balancer, providing a corrective control to mitigate denial-of-service attacks.
- Sensors attached to mirrored switch ports enable intrusion detection for the most sensitive privilege level hosts or zones.

On each host, endpoint protection software applies a range of preventive, detective, and corrective controls to mitigate threats that have evaded network controls. Endpoint software can implement host firewalls, anti-virus, intrusion detection, and data loss prevention.

Device Attributes

Attributes determine the precise way in which a device can be placed within the network topology.

Active versus Passive

A passive security control is one that does not require any sort of client or agent configuration or host data transfer to operate. For example, network traffic can be directed or copied to a sensor and scanned by an analysis engine. This control is completely passive. Hosts on the network would be unaware that it is operating. The control has no addressable interface.

An active security control that performs scanning must be configured with credentials and access permissions and exchange data with target hosts. An active control that performs filtering requires hosts to be explicitly configured to use the control. This might mean installing agent software on the host or configuring network settings to use the control as a gateway.

Inline versus Tap/Monitor

A device that is deployed inline becomes part of the cable path. No changes in the IP or routing topology are required. The device's interfaces are not configured with MAC or IP addresses.

As an example of inline versus monitored deployment options, controls that sniff network traffic can be deployed via a sensor attached to a switch or via a tap attached to a network cable:

- SPAN (switched port analyzer)/mirror port —this means that the sensor is attached to a specially configured port on a switch that receives copies of frames addressed to nominated access ports (or all the other ports). This method is not completely reliable. Frames with errors will not be mirrored , and frames may be dropped under heavy load.
- Test access point (TAP) —this is a box with ports for incoming and outgoing network cabling and an inductor or optical splitter that physically copies the signal from the cabling to a monitor port. There are types for copper and fiber optic cabling. Unlike a SPAN, no logic decisions are made so the monitor port receives every frame—corrupt or malformed or not—and the copying is unaffected by load.

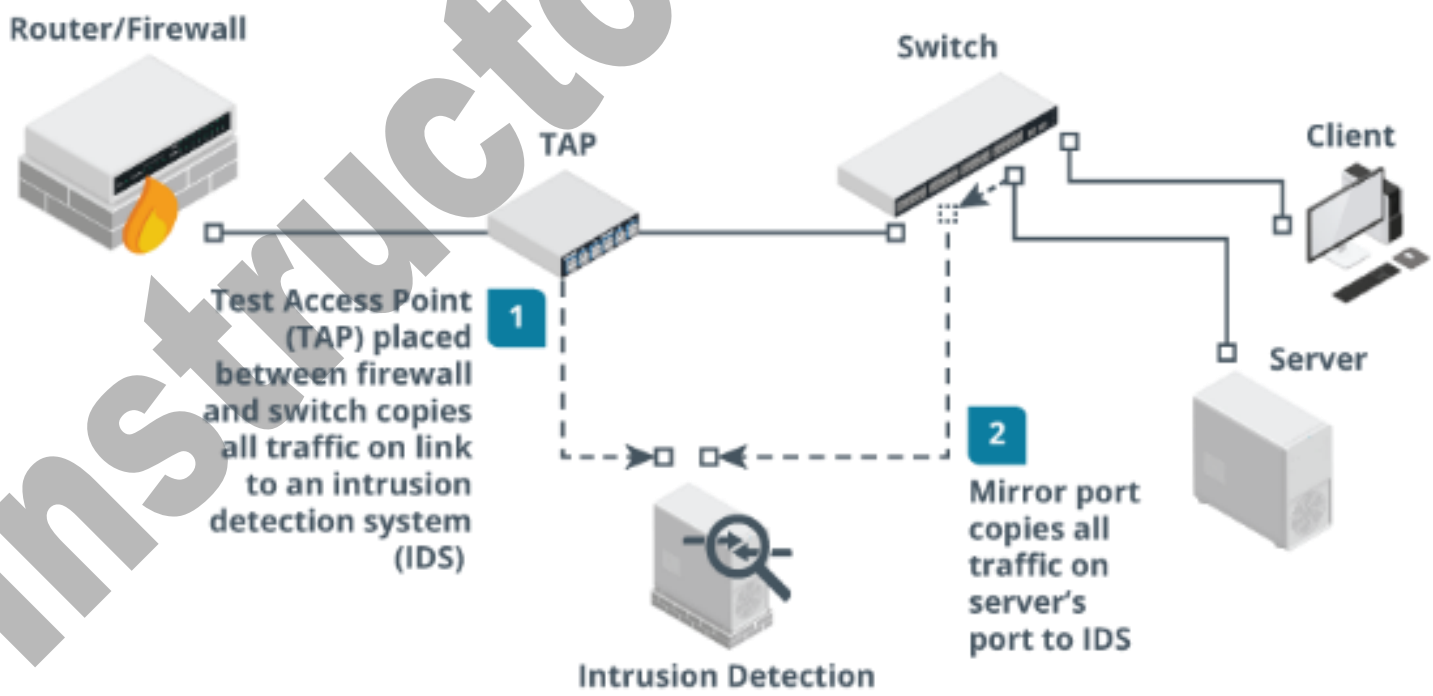

A TAP device is placed inline with the cable path, while a mirror port uses the switch to copy frames to a detection system. The router/firewall is an active control as client devices must be configured to use it for Internet access. The TAP and mirror ports are passive controls. The are completely transparent to the server and client hosts.

A security device could enter a failure state for a number of reasons. There could be a power or hardware fault, an irreconcilable policy violation, or a configuration error. Hardware failure can be caused by power surges, overheating, and physical damage. Software failure can occur because of bugs, security vulnerabilities, and compatibility issues. Configuration issues can be caused by human errors such as inattention, fatigue, or lack of training. Finally, devices or sites might be impacted by natural disasters such as floods, hurricanes, and earthquakes.

When it fails, a device can designed or configured to fail-open or fail-closed:

- Fail-open means that network or host access is preserved, if possible. This mode prioritizes availability over confidentiality and integrity. The risk of a fail-open control is that a threat actor could engineer a failure state to defeat the control.
- Fail-closed means that access is blocked or that the system enters the most secure state available, given whatever failure occurred. This mode prioritizes confidentiality and integrity over availability. The risk of a fail-closed control is system downtime.

It may or may not be possible to configure the fail mode. For example, an inline security appliance that suffers power failure will fail-closed unless there is an alternative network path. Some devices designed to be installed inline have a backup cable path that will allow a fail-open operation.

5.1.3 Practice Questions (Section Quiz)

q_ent_architecture_considerations_secp8

You are a network architect for a rapidly growing startup. The startup is planning to expand its operations and is considering a major upgrade to its network architecture.

Which of the following factors should be your primary consideration when designing the new network architecture?

Answers:

- Minimizing initial capital outlay by choosing the cheapest available hardware and software options.
- Maximizing compute resources and responsiveness, regardless of cost.
- ***Balancing costs, compute and responsiveness, scalability, availability, and resilience.**
- Prioritizing scalability and ease of deployment over all other considerations.

Explanation:

Balancing costs, compute and responsiveness, scalability, availability, and resilience is the best approach to network architecture design. This approach ensures that the network can meet the startup's current and future needs, while also considering cost-effectiveness and the ability to recover from potential failures.

Minimizing initial capital outlay is not the best choice because while minimizing initial capital outlay is important, choosing the cheapest available options could lead to performance issues, lack of scalability, and potential security risks. The long-term costs of such a decision could outweigh the initial savings.

Maximizing compute resources and responsiveness is not the best choice because while maximizing compute resources and responsiveness is important, doing so regardless of cost could lead to unnecessary expenditure. It's important to balance these factors with cost considerations.

Prioritizing scalability and ease of deployment is not the best choice because while scalability and ease of deployment are important considerations, they should not be prioritized over all other considerations. Factors such as cost, compute and responsiveness, availability, and resilience are also important and should be balanced with scalability and ease of deployment.

q_ent_architecture_defense_in_depth_secp8

You are a cybersecurity specialist for a financial institution that is planning to enhance its network security. The institution has decided to adopt a defense in depth strategy.

Which of the following approaches would BEST align with a defense in-depth strategy?

Answers:

- Implementing a single, robust firewall at the network perimeter and relying on this for all network security.
- ***Implementing multiple security measures at different network layers, including firewalls, intrusion detection systems, and regular patch management.**
- Implementing a single security measure, such as encryption, across all data and network traffic.
- Implementing a complex array of different security technologies without considering their interaction or potential redundancies.

Explanation:

Implementing multiple security measures at different network layers aligns with the defense in-depth strategy is the correct answer. This approach ensures that if one security measure fails, others are in place to provide protection.

While a robust firewall is an important part of network security, relying on a single security measure does not align with the defense in-depth strategy, which advocates for multiple layers of security.

While encryption is a valuable security measure, relying solely on it does not provide the multiple layers of security advocated for in a defense in-depth strategy.

While a defense in-depth strategy does involve implementing multiple security measures, these measures should be carefully considered and coordinated. Implementing a complex array of different security technologies without considering their interaction or potential redundancies could lead to gaps in security and inefficient use of resources.

q_ent_architecture_fail-closed_01_secp8

A hospital has implemented a security device that processes sensitive patient information. The hospital wants to ensure that in the event of a failure, the confidentiality and integrity of the patient data take priority over the system's availability.

What should the hospital set as the failure mode configuration for this security device?

Answers:

- ***The security device should be configured to fail-closed.**
- The security device should be configured to fail-open.
- The security device should be configured to actively monitor the network.
- The security device should be configured to passively monitor the network.

Explanation:

A fail-closed configuration prioritizes confidentiality and integrity over availability. In the event of a failure, a fail-closed device would block access or enter the most secure state available, protecting patient data.

A fail-open configuration would prioritize availability over confidentiality and integrity. In the event of a failure, a fail-open device would preserve network or host access, potentially exposing sensitive patient data.

While active monitoring is an important part of network security, it does not directly pertain to the failure mode of the device and would not ensure the confidentiality and integrity of patient data.

As with active monitoring, passive monitoring is important for network security, but it does not directly pertain to the failure mode of the device.

q_ent_architecture_fail-closed_02_secp8

A financial institution is implementing a new security control device to protect its network infrastructure and wants to ensure that in the event of a failure, the confidentiality and integrity of its financial data take precedence over system availability.

What should the financial institution set as the failure mode configuration for this security control device?

Answers:

- ***The security control device should be configured to fail-closed.**
- The security control device should be configured to fail-open.
- The security control device should be configured to actively monitor the network.
- The security control device should be configured to passively monitor the network.

Explanation:

Configuring the security control device to fail-closed ensures that confidentiality and integrity take precedence over availability, protecting financial data during a failure.

Fail-open configuration prioritizes availability over confidentiality/integrity and may expose sensitive financial data during a failure by preserving network/host access.

While active monitoring is important, it does not directly address the required failure mode configuration to prioritize confidentiality and integrity in this scenario.

Passive monitoring, although useful for network security, does not directly pertain to the failure mode configuration required to safeguard the confidentiality and integrity of financial data during failures.

q_ent_architecture_fail-open_01_secp8

An organization implements a new network infrastructure and plans to use an intrusion prevention system (IPS) for security. The IT manager wants to ensure that the IPS will continue to let traffic flow if it fails.

Which failure mode should the IT manager configure the IPS?

Answers:

- ***Fail-open**
- Fail-closed
- Active

- Passive

Explanation:

In a fail-open mode, if the IPS fails, it will still allow traffic to pass through, maintaining network connectivity

In a fail-closed mode, if the IPS fails, it will block all traffic, disrupting network connectivity

Active mode is an IPS operation mode, not a failure mode. In active mode, the IPS monitors traffic and blocks threats. A fail-open mode ensures traffic flow if the IPS fails.

Passive mode is an IPS operation mode, not a failure mode. In passive mode, the IPS monitors traffic and generates alerts but doesn't block threats.

q_ent_architecture_fail-open_02_secp8

An organization is implementing an Intrusion Prevention System (IPS) as part of its efforts to secure its enterprise infrastructure. The IT manager is considering the failure modes of the IPS and is deciding between a fail-open and a fail-closed configuration.

What are the implications of each configuration on network traffic in the event of a system failure?

Answers:

- ***Fail-open will allow all traffic, while fail-closed will block all traffic.**
- Fail-open will block all traffic, while fail-closed will allow all traffic.
- Both fail-open and fail-closed will block all traffic.
- Both fail-open and fail-closed will allow all traffic.

Explanation:

An IPS can either allow all traffic through without disruption but with a risk of malware (fail-open) or block all traffic without malware but with a risk of disruption (fail-closed).

In the event of a system failure, a fail-open configuration allows all traffic, while a fail-closed configuration blocks all traffic.

Even if the system fails, a fail-open configuration allows all traffic to pass through, ensuring that there is no disruption to the flow of data and information.

If the system fails, a fail-closed configuration blocks all traffic from passing through, ensuring that no malware or other harmful content can enter the network.

q_ent_architecture_inline_mode_secp8

The IT manager of a medium-sized organization is designing a new network infrastructure to secure its enterprise infrastructure by implementing an Intrusion Prevention System (IPS) and an Intrusion Detection System (IDS). The manager is considering different deployment methods for the IPS/IDS to optimize their effectiveness.

The organization's network includes multiple security zones, a virtual private network (VPN) for remote access, and a web application firewall (WAF).

Which deployment method provides the MOST comprehensive protection in this scenario?

Answers:

- ***Deploy the IPS/IDS devices in inline mode at the network perimeter.**
- Deploy the IPS/IDS devices in passive mode within the internal network.
- Deploy the IPS/IDS devices in tap/monitor mode at the entry and exit points of the VPN tunnel.
- Deploy the IPS/IDS devices in inline mode next to the WAF.

Explanation:

Deploying the IPS/IDS devices in inline mode at the network perimeter allows for real-time analysis and reaction to potential threats, providing comprehensive protection for all inbound and outbound network traffic.

While deploying the IPS/IDS devices in passive mode inside the internal network provides visibility into internal threats, it does not provide real-time threat mitigation or cover inbound and outbound traffic.

IPS/IDS devices in tap/monitor mode at the entry and exit points of the VPN tunnel can help monitor remote access but do not protect all network traffic.

Although deploying the IPS/IDS devices in inline mode next to the WAF can enhance the security of web applications, it does not protect all network traffic.

q_ent_architecture_load_balancer_secp8

A small start-up has recently launched its first web application. To ensure high availability and to handle potential traffic spikes, the start-up decides to implement a load balancer in its network infrastructure.

The network technician must secure the load balancer against basic threats.

What is the fundamental step the network technician should take to secure the load balancer?

Answers:

- ***Disable unnecessary services on the load balancer.**
- Implement an intrusion detection system (IDS) alongside the load balancer.
- Configure the load balancer to operate in fail-closed mode.
- Enable all available features on the load balancer.

Explanation:

Disabling unnecessary services on the load balancer is a fundamental step in reducing the attack surface and enhancing security.

While implementing an IDS can enhance security, it is not a fundamental step when securing a load balancer and might be an overkill for a small start-up.

Configuring the load balancer to operate in a fail-closed mode can prevent unfiltered network traffic in case of device failure. However, it is not the most basic step in securing a load balancer.

Enabling all available features on the load balancer is not a good security practice. It could potentially open up unnecessary vulnerabilities and complicate the configuration.

q_ent_architecture_network_nodes_secp8

You are a network engineer for a multinational corporation. The corporation is planning to expand its operations to a new location and you are tasked with designing the network for the new site. The network should be robust, scalable, and secure.

Which of the following approaches to setting up network nodes at the new site would best meet these requirements?

Answers:

- Setting up a single powerful server to handle all network functions.
- ***Setting up multiple network nodes, each dedicated to a specific function such as routing, switching, and firewalling.**
- Setting up a single network node with minimal functionality to reduce costs.
- Setting up multiple network nodes with redundant functionality to ensure high availability.

Explanation:

Setting up multiple network nodes, each dedicated to a specific function, is the correct answer and allows for better performance, scalability, and security. Each node can be optimized for its specific function, and if one node fails, the impact on the overall network is minimized.

Setting up a single power server is not the best choice because a single server handling all network functions could become a bottleneck, affecting network performance. Additionally, if this server fails, all network functions would be affected, leading to significant downtime.

Setting up a single network node is not the best choice because while it may reduce initial costs, a single network node with minimal functionality would likely not be able to handle the network demands of a multinational corporation. This could lead to poor network performance and potential security risks.

Setting up multiple network nodes with redundant functionality is not the best choice because while redundancy is important for high availability, having multiple network nodes with redundant functionality could lead to unnecessary complexity and costs. It would be more efficient to have each network node dedicated to a specific function, with appropriate backup and failover mechanisms in place.

q_ent_architecture_network_protocols_secp8

You are a network engineer for a global company that is implementing a new real-time data processing system. This system requires efficient and reliable data transfer between different network segments.

Which of the following network components would be most critical in ensuring the efficient and reliable transfer of real-time data in this scenario?

Answers:

- Application protocols
- Domain Name System (DNS) servers
- ***Transport protocols**
- Wireless access points

Explanation:

Transport protocols is the correct answer. Transport protocols, such as the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP), operate at the transport layer (Layer 4) of the OSI model. They are responsible for the end-to-end transfer of data between systems and can provide mechanisms for ensuring reliable data transfer, which is critical for real-time data processing systems.

Application protocols are not the best choice because they operate at the application layer (Layer 7) of the OSI model and are more concerned with the format and control of the data exchange between applications, not the efficient and reliable transfer of data between network segments.

Domain Name System (DNS) servers are not the best choice because their primary function is to translate domain names into IP addresses, not to ensure the efficient and reliable transfer of real-time data between network segments.

Wireless access points are not the best choice because their primary function is to provide a bridge between a wired network and wireless devices. While they do play a role in data transfer, they do not have the same level of control over the efficiency and reliability of data transfer as transport protocols.

q_ent_architecture_osi_secp8

You are a network architect for a large organization. The organization is planning to upgrade its network infrastructure to support a new business application. The application requires high availability, secure data transfer, and efficient handling of large data volumes.

Which of the following network design considerations best aligns with the requirements of the new business application?

Answers:

- Implementing a flat network design with all devices on the same network segment.
- ***Implementing a layered network design based on the OSI model with appropriate security controls at each layer.**
- Implementing a network design with minimal use of routers to reduce complexity.
- Implementing a network design that prioritizes cost reduction over scalability and resilience.

Explanation:

A layered network design based on the OSI model is the correct answer. The OSI model allows for the implementation of appropriate security controls at each layer, from the physical layer up to the application layer. This design also supports high availability and efficient handling of large data volumes by allowing for network segmentation and the use of specialized devices and protocols at each layer.

A flat network design is not the best choice because a flat network design does not provide the necessary segmentation for secure data transfer and efficient handling of large data volumes. In a flat network, all devices are on the same network segment, which can lead to performance issues and increased security risks.

Minimal use of routers is not the best choice because, while it may reduce complexity, it does not support high availability or efficient handling of large data volumes. Routers play a crucial role in directing network traffic and segmenting the network, which are important for the requirements of the new business application.

Prioritizing cost reduction over scalability and resilience is not the best choice because while cost reduction is an important consideration, it should not be prioritized over scalability and resilience, especially for a business application that requires high availability and secure data transfer. A network design that lacks scalability and resilience may result in higher costs in the long run due to potential downtime and security incidents.

q_ent_architecture_risk_transference_secp8

As a senior network architect at your company, you have been tasked with evaluating the current network architecture and proposing changes to improve security and efficiency.

The company has been experiencing frequent network outages due to an outdated infrastructure, which has led to significant financial losses. The board of directors is concerned about the costs associated with a complete network overhaul and is considering other options.

Which of the following is the MOST appropriate course of action?

Answers:

- Implement a complete network overhaul immediately, regardless of the costs.
- Do nothing and continue with the current network architecture.
- ***Transfer the risk by outsourcing the network management to a third-party service provider.**
- Invest in minor upgrades to the current network infrastructure.

Explanation:

Transferring the risk by outsourcing the network management to a third-party service provider is the correct answer. This approach allows the company to leverage the expertise of a third-party provider, potentially improving network stability and security. It also transfers the risk associated with network management to the third-party provider, aligning with the concept of risk transference.

Implementing a complete network overhaul immediately, regardless of the costs, may not be the best solution. While it could potentially solve the current issues, it may also introduce new ones and could be financially burdensome for the company. This approach does not consider the risk transference option.

Doing nothing and continuing with the current network architecture is not a viable option. The frequent network outages are causing significant financial losses and damaging the company's reputation. This approach ignores the problem and does not consider risk transference.

Investing in minor upgrades to the current network infrastructure may not be sufficient to address the frequent network outages. While this approach may improve the situation in the short term, it does not provide a long-term solution and does not consider risk transference.

q_ent_architecture_tap_secp8

You are a network engineer for a large corporation that is planning to implement a new intrusion detection system (IDS). The corporation has a high volume of network traffic and requires real-time monitoring for potential security threats.

Which of the following approaches to integrating the IDS into the corporation's network would best meet these requirements?

Answers:

- Connecting the IDS directly to the corporation's main router.
- ***Using a Test Access Point (TAP) device to provide the IDS with a copy of the network traffic.**
- Connecting the IDS to a switch and configuring the switch to mirror all network traffic to the IDS.
- Connecting the IDS to a random network node to minimize impact on network performance.

Explanation:

Using a Test Access Point (TAP) device is the correct answer. A TAP device allows the IDS to receive a copy of all network traffic without impacting network performance. This approach enables real-time monitoring and is suitable for networks with high traffic volumes.

Connecting the IDS directly to the corporation's main router is not the best choice because connecting the IDS directly to the corporation's main router could impact network performance. The router may not be able to handle the additional load of forwarding all network traffic to the IDS.

Connecting the IDS to a switch and configuring the switch to mirror all network traffic to the IDS is not the best choice because while a switch can be configured to mirror all network traffic to the IDS, this approach could impact the switch's performance and may not be suitable for networks with high traffic volumes.

Connecting the IDS to a random network node to minimize impact on network performance is not the best choice because connecting the IDS to a random network node would not ensure that the IDS receives all network traffic. This approach could result in incomplete monitoring and potential security risks.

5.2 Security Appliances

As you study this section, answer the following questions:

- What are the benefits and risks of using proxy servers?
- What is the purpose of a content filtering server?
- What are the uses of a screened subnet?
- Why is a honeynet useful?
- What are the features of an all-in-one security appliance?
- What size organization should employ a all-in-one security appliance?

In this section, you will learn to:

- Configure a security appliance.
- Configure network security appliance access.

The key terms for this section include:

Term	Definition
Security zone	Portions of the network or system that have specific security concerns or requirements.
Wireless network	A network that does not require a physical connection.
Guest network	A network that grants internet access only to guest users. A guest network has a firewall to regulate guest user access.
Honeynet	A host (honeypot), network (honeynet), file (honeyfile), or credential/token (honeytoken) set up with the purpose of luring attackers away from assets of actual value and/or discovering attack strategies and weaknesses in the security configuration.
Ad hoc	A decentralized network that allows connections without a traditional base station or router. It allows users to connect two or more devices directly to each other for a specific purpose.

DNS sinkhole	A temporary DNS record that redirects malicious traffic to a controlled IP address.
Jump server	A hardened server that provides access to other hosts.
Agent-based filtering	Agent-based web filtering involves installing a software agent on desktop computers, laptops, and mobile devices. The agents enforce compliance with the organization's web filtering policies.
Screened subnet	A network that contains publicly accessible resources and is located between the private network and an untrusted network, such as the internet. It is protected by a firewall.
Proxy server	A type of firewall that stands as an intermediary between clients requesting resources from other servers.
Internet content filter	Software used to monitor and restrict content delivered across the web to an end user.
Fake telemetry	Deception strategy that returns spoofed data in response to network probes.
All-in-one security appliance	An appliance that combines many security functions into a single device.
Application-aware devices	A device that has the ability to analyze and manage network traffic based on the application-layer protocol.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts.
	• Deception and disruption technology
	  ○ Honeypot
	  ○ Honeynet
	  ○ Honeyfile
	  ○ Honeytoken
	3.2 Given a scenario, apply security principles to secure enterprise infrastructure.
	• Infrastructure considerations
	  ○ Security zones
	• Failure modes

	○ Fail-open ○ Fail-closed • Network appliances ○ Jump server ○ Proxy server ○ Load balancer ○ Sensors 4.5 Given a scenario, modify enterprise capabilities to enhance security. • Web filter ○ Agent-based ○ Centralized proxy ○ Universal Resource Locator (URL) scanning ○ Content categorization ○ Block rules ○ Reputation
TestOut Security Pro	2.1 Harden Physical Access 2.1.2 Install and Configure a Security Appliance 2.1.4 Create and Configure a screened subnet

5.2.1 Security Appliances (Lesson Video)

Transcript:

In network security, many appliances are available to help secure your network. The size, needs, and resources of your organization will be factors in determining which ones to implement for your network. In this lesson, we'll look at a few of these options.

Let's start with proxy servers. A proxy is a person who has the authority to act on behalf of another person.
Similarly, a proxy server acts on behalf of a client or user when attempting to access resources over the internet. The proxy server, in this position of an intermediary, provides a layer of protection to the client. A proxy server works on a store-and-forward model. This means the proxy deconstructs each packet, performs analysis, then rebuilds the packet and forwards it on if it conforms to the rules it's been configured with.
The main benefit of a proxy is that client computers connect to a specified point on the perimeter network for web access. This provides a degree of traffic management and security. Proxy servers can also provide anonymity for users by masking their IP addresses. In addition, most web proxy servers provide caching engines; the proxy retains frequently requested web pages, eliminating the need to re-fetch those pages for subsequent requests.
Proxy servers are also often used for content filtering and monitoring. They can block certain websites or filter out malicious content before it reaches the client. Proxy servers can even distribute incoming network traffic across multiple servers, helping to balance the load and improve the overall performance and reliability of a network.
The next appliance is a jump server. A jump server is a hardened server that provides access to other hosts. Jump servers are often used for administrative tasks, where administrators connect to the jump server first and then use it to access other internal systems, like servers or devices.
A jump server is primarily used to enhance security by controlling access to sensitive resources. It acts as a gateway to access certain systems kept isolated from the external network. The jump server is typically locked down and secured to a higher degree, ensuring that only authorized users can access it, reducing the risk of unauthorized access to critical systems.
Now, let's talk about load balancers. While primarily used to distribute network traffic across multiple servers to optimize performance, a load balancer can also serve as a security appliance in certain scenarios. A load balancer distributes

client requests across available server nodes in a farm or pool. This is used to provision services that can scale from light to heavy loads and to provide mitigation against denial-of-service attacks.
A load balancer also provides fault tolerance. When there are several servers in a farm, all identified by a single name or IP address via a load balancer, if one server stops working, client requests can be sent to another server in the same group.
A load balancer can be deployed in any situation where there are multiple servers providing the same function. Examples include web servers, front-end email servers, web conferencing, and streaming media servers.
There are two main types of load balancers: layer 4 and layer 7. A layer 4 load balancer makes forwarding decisions on IP address and TCP/UPD port values, working at the transport layer of the OSI model. A layer 7 load balancer, or content switch, makes forwarding decisions based on application-level data, such as a request for a particular URL web address or data types like video or audio streaming, which requires more complex logic.
Load balancers can also include a Web Application Firewall that provides protection against common web application attacks. Modern load balancers often include features for traffic inspection and filtering. They can analyze incoming traffic for suspicious patterns, such as unusual request rates or malicious payloads. Traffic from potentially harmful sources can be dropped or redirected for further analysis.
Another network security appliance is a packet sniffer, which is referred to as a sensor. Typically, the packet capture sensor is placed behind a firewall or close to a server of particular importance. The idea is to identify malicious traffic that has managed to get past the firewall. A single sniffer can record a large amount of traffic data, so don't put multiple sensors all over the network without provisioning the resources to manage them properly. Depending on network size and resources, one or only a few sensors are typically deployed to monitor key assets or network paths.
The traffic captured by each sensor is transferred to a host or appliance running an intrusion detection system, or IDS, such as Snort, Suricata, or Zeek. When traffic matches a detection signature, the IDS raises an alert or generates a log entry but does not block the source host. This type of passive sensor doesn't slow down traffic and is undetectable by the attacker.
An IDS is used to identify and log hosts and applications and to detect password-guessing attempts, port scans, worms, backdoor applications, malformed packets or sessions, and other policy violations.
The last appliance we'll look at in this lesson is an all-in-one security appliance. An all-in-one security appliance incorporates multiple security functions into a single piece of hardware.
Unified Threat Management, or UTM, is the most common all-in-one appliance. UTM puts several key security components into a single device, usually managed using a web interface.
Let's look at some of the functionality commonly implemented in all-in-one security devices. One of the first features is URL filtering, which prevents users from accessing URL categories restricted by the organization. These often include pornographic, shopping, and gambling sites.
Another feature is content inspection. This helps to ensure that HTTP connections and content meet certain criteria. For example, many content filters actively monitor data streams by inspecting the packets in search of viruses, Trojans, worms, and other malicious code.
These devices also include a spam filter to reduce the chance that your users' mailboxes will fill up with junk email. The UTM will also probably include a firewall so you can set up rules and log traffic to and from the network. Many devices include an integrated intrusion detection or prevention system that alerts you if a network attack attempt is detected or prevented.
An all-in-one device typically contains networking features as well as UTM features. These network features include an integrated switch. For small and remote offices that only have a handful of systems, a small integrated switch provides all the connectivity required. As part of switching, the device may include a network router that allows you to connect to a network backbone using an uplink port.
The device may also support traffic shaping to prioritize one type of traffic over another. For example, if the organization is using voice-over IP phones, the all-in-one appliance can give voice traffic priority over normal data traffic.
And that's it for this lesson. In this lesson, we discussed a few security appliance options. First, we talked about proxy servers and how they work as middlemen between clients and internet resources. Then, we looked at jump servers that are used to control access to sensitive resources. Then, we reviewed load balancers and what role they can play in security. After that, we looked at sensors and how they record traffic data and supply that information to an IDS, which flags anything with a matching detection signature. We ended this lesson by discussing an all-in-one security appliance called a Unified Threat Management appliance, which can combine multiple devices into one appliance.

5.2.2 Security Solutions (Lesson Video)

Transcript:

Many of us take for granted the clean water we get from our faucets every day. Before we take a drink, the dangerous contaminants have been filtered out, leaving us with only what we want. One very important security solution for networks is filtering: content filtering and web filtering. Just like filtering out dangerous contaminants from our drinking water, content filters and web filters filter out dangerous contaminants to our networks and users. Let's take a closer look.

Let's start with content filtering servers. Content filtering servers typically have one or two configurations. Either you allow all content except for what you decide to ban, or you block all content except for what you decide to allow. These controls are implemented through security levels, allow lists, deny lists, and category levels.

A content filter is designed to give network administrators and management the ability to block unwanted or inappropriate website browsing. For example, several organizations have guidelines stating that users can't visit adult sites, gambling, or shopping sites while at work. A content filter can be put in place, and the network administrator can configure the filter to prevent users from visiting these types of sites.

Most often, the content filter allows everything except sites and categories that are blocked via a denylist. For additional flexibility, security groups can be configured to allow or prevent site access by group.

New websites are created every day, so the filters must be updated regularly, like virus definitions. Most content filtering servers have an online subscription service to update websites and website categories at certain intervals, usually daily or weekly.

Now, web filtering plays a pivotal role in safeguarding an organization's network. Its primary function is to block users from accessing malicious or inappropriate websites, thereby protecting the network from potential threats.

Web filters analyze web traffic, often in real-time, and can restrict access based on various criteria such as URL, IP address, content category, or even specific keywords.

Now, let's look at two different approaches to web filtering: agent-based and centralized. The first approach is focused on filtering at each endpoint. The second filters from a centralized proxy server that endpoints are connected to.

Agent-based web filtering involves installing a software agent on desktop computers, laptops, and mobile devices. The agents enforce compliance with the organization's web filtering policies. Agents communicate with a centralized management server to retrieve filtering policies and rules and then apply them locally on the device. Agent-based solutions typically leverage cloud platforms to ensure they can communicate with devices regardless of the network they're connected to. This means filtering policies remain in effect even when users are off the corporate network, such as when working from home or traveling.

Because filtering occurs locally on the device, agent-based methods often provide more granular control, such as filtering HTTPS traffic or applying different filtering rules for different applications. Agent-based filtering can also provide detailed reporting and analytics.

In a centralized approach, a proxy server performs content filtering by analyzing web requests from users and determining whether to permit or deny access based on established policies. A centralized proxy server acts as an intermediary between end users and the internet. Beyond blocking unwanted or harmful content, a centralized proxy server can also perform detailed logging and reporting of web activity.

A centralized proxy server employs various techniques to protect web traffic and ensure the safety of an organization's network. A proxy server can perform URL scanning, which involves blocking URLs known to host malicious content, are inappropriate, or violate the company's policies.

Proxy servers also allow content categorization. This means that websites can be classified into various categories like social networking, gambling, webmail, and many others. Organizations can define rules to allow or deny access based on these categories.

Another technique proxy servers can implement is block rules. For example, an organization could block all .exe downloads or certain domains or content categories.

Proxy servers can even incorporate reputation-based filtering, which leverages continually updated databases that score websites based on their observed behavior and history. Sites known for hosting malware, engaging in phishing attacks, or distributing spam, for instance, would have a poor reputation score and could be automatically blocked.

Content and web filtering aren't without potential issues and challenges. One common problem is overblocking or underblocking. Overblocking occurs when the filter is too restrictive, inadvertently blocking access to legitimate and useful websites and negatively impacting employee productivity.

Underblocking, on the other hand, occurs when the filter allows access to potentially harmful or inappropriate websites.

Another issue is the handling of encrypted traffic like HTTPS. Without proper configuration, web filters may be unable to inspect encrypted traffic, representing most modern web traffic.

Well, that's it for this lesson. In this lesson, we discussed security solutions based on content and web filtering. We looked at how content filters can either allow everything except what you specifically ban or they can ban everything

except what's specifically allowed. We then looked at two approaches to web filtering: agent-based, which is applied to endpoints, and centralized, which is provided through a proxy server that network traffic flows through. We ended the lesson by looking at the challenge of not overblocking or underblocking when applying filters so that the network is protected, and users can do their work and access their resources.

5.2.3 Security Solution Facts

This lesson covers the following topics:

- Security appliances
- Security solutions

Security Appliances

In network security, there are many appliances available to help secure your network. The size, needs, and resources of your organization will be factors in determining which ones to implement for your network. The following table includes several security appliances.

Appliance	Description
Proxy server	A proxy server acts on behalf of a client or user when attempting to access resources over the internet. The proxy server, in this position of an intermediary, provides a layer of protection to the client. A proxy server works on a store-and-forward model. This means the proxy deconstructs each packet, performs analysis, then rebuilds the packet and forwards it on, if it conforms to the rules it's been configured with. Client computers connect to a specified point on the perimeter network for web access. Benefits a proxy server can provide: • Traffic management • Protection • Anonymity for users by masking their IP addresses • Caching engines • Content filtering • Content monitoring • Incoming network traffic distribution across multiple servers to help balance the load
Jump server	A jump server is a hardened server that provides access to other hosts. Jump servers are often used for administrative tasks where administrators connect to the jump server first and then use it to access other internal systems, like servers or devices. A jump server is primarily used to enhance security by controlling access to sensitive resources. It acts as a gateway to access certain systems that are kept isolated from the external network. The jump server is typically locked down and secured to a higher degree, ensuring that only authorized users can access it, reducing the risk of unauthorized access to critical systems.
Load balancer	A load balancer, while primarily used to distribute network traffic across multiple servers to optimize performance, can also serve as a security appliance in certain scenarios. A load balancer distributes client requests across available server nodes in a farm or pool. This is used to provision services that can scale from light to heavy loads and to provide mitigation against denial-of-service attacks. A load balancer also provides fault tolerance. If there are multiple servers available in a farm all addressed by a single name/IP address via a load balancer, then if a single server fails, client

Appliance	Description
	requests can be forwarded to another server in the farm. A load balancer can be deployed in any situation where there are multiple servers providing the same function. Examples include web servers, front-end email servers, web conferencing, video conferencing, and streaming media servers. There are two main types of load balancers: • Layer 4 — A layer 4 load balancer makes forwarding decisions on IP address and TCP/UPD port values, working at the transport layer of the OSI model. • Layer 7 — A layer 7 load balancer, or content switch, makes forwarding decisions based on application-level data, such as a request for a particular URL web address or data types like video or audio streaming, which requires more complex logic. Load balancers can also include a Web Application Firewall (WAF) that provides protection against common web application attacks. Modern load balancers often include features for traffic inspection and filtering. They can analyze incoming traffic for suspicious patterns, such as unusual request rates or malicious payloads. Traffic from potentially harmful sources can be dropped or redirected for further analysis.
Sensor	A packet sniffer is referred to as a sensor. Typically, the packet capture sensor is placed behind a firewall or close to a server of particular importance. The idea is to identify malicious traffic that has managed to get past the firewall. A single sniffer can record a large amount of traffic data so it's best to not put multiple sensors all over the network without provisioning the resources to manage them properly. Depending on network size and resources, one or only a few sensors are deployed to monitor key assets or network paths. The traffic captured by each sensor is transferred to a host or appliance running an intrusion detection system (IDS), such as Snort, Suricata, or Zeek. When traffic matches a detection signature, the IDS raises an alert or generates a log entry but does not block the source host. This type of passive sensor does not slow down traffic and is undetectable by the attacker. An IDS is used to identify and log hosts and applications and to detect password-guessing attempts, port scans, worms, backdoor applications, malformed packets or sessions, and other policy violations.
All-in-one security appliance	An all-in-one security appliance incorporates multiple security functions into a single piece of hardware. Unified Threat Management, or UTM, is the most common all-in-one appliance. UTM puts several key security components into a single device that's usually managed using a web interface. A manufacturer subscription for updates may also be required. Commonly implemented functionality of all-in-one-security devices include the following: • URL filtering—prevents users from accessing URL restricted categories. • Content inspection—helps ensure HTTP connections and content meet specified criteria. For example, many content filters actively monitor data streams by inspecting the packets in search of viruses, Trojans, worms, and other malicious code. • Spam filtering—reduces junk mail in mailboxes of users. • Firewall—can be configured with rules and used to log traffic to and from the network. • Intrusion detection system—An IDS sends an alert if a network attack is detected. • Intrusion prevention system—An IPS sends an alert if a network attack is detected or prevented.

Appliance	Description
	A UTM usually includes networking features as well, such as the following: • Switch • Router • Traffic shaping management

Failure Modes

A security device could enter a failure state for a number of reasons. There could be a power or hardware fault, an irreconcilable policy violation, or a configuration error. Hardware failure can be caused by power surges, overheating, and physical damage. Software failure can occur because of bugs, security vulnerabilities, and compatibility issues. Configuration issues can be caused by human errors such as inattention, fatigue, or lack of training. Finally, devices or sites might be impacted by natural disasters such as floods, hurricanes, and earthquakes.

When it fails, a device can designed or configured to fail-open or fail-closed:

- Fail-open means that network or host access is preserved, if possible. This mode prioritizes availability over confidentiality and integrity. The risk of a fail-open control is that a threat actor could engineer a failure state to defeat the control.
- Fail-closed means that access is blocked or that the system enters the most secure state available, given whatever failure occurred. This mode prioritizes confidentiality and integrity over availability. The risk of a fail-closed control is system downtime.

It may or may not be possible to configure the fail mode. For example, an inline security appliance that suffers power failure will fail-closed unless there is an alternative network path. Some devices designed to be installed inline have a backup cable path that will allow a fail-open operation.

Security Solutions

Content Filtering Server

A content filtering server can deny or allow users access to certain websites. Content filtering servers typically have one or two configurations:

- Allow all content except for what you decide to ban
- Block all content except for what you decided to allow.

These controls are implemented through security levels, allow lists, deny lists, and category levels.

A content filter is designed to give network administrators and management the ability to block unwanted or inappropriate website browsing. For example, several organizations have guidelines stating that users can't visit adult sites, gambling, or shopping sites while at work. A content filter can be put in place, and the network administrator can configure the filter to prevent users from visiting these types of sites.
Most often, the content filter allows everything except sites and categories that are blocked via a deny list. For additional flexibility, security groups can be configured to allow or prevent site access by group.
New websites are created every day, so the filters must be updated regularly, like virus definitions. Most content filtering

servers have an online subscription service to update websites and website categories at certain intervals, usually daily or weekly.

Web Filtering

Web filtering plays a pivotal role in safeguarding an organization's network. Its primary function is to block users from accessing malicious or inappropriate websites, thereby protecting the network from potential threats. Web filters analyze web traffic, often in real time, and can restrict access based on various criteria such as URL, IP address, content category, or even specific key words.
Two approaches to web filtering include:

- Agent-based which is focused on filtering at each endpoint.
- Centralized filters from a centralized proxy server to which endpoints are connected.

Agent-Based Filtering

Agent-based web filtering involves installing a software agent on desktop computers, laptops, and mobile devices. The agents enforce compliance with the organization's web filtering policies. Agents communicate with a centralized management server to retrieve filtering policies and rules and then apply them locally on the device. Agent-based solutions typically leverage cloud platforms to ensure they can communicate with devices regardless of the network they are connected to. This means filtering policies remain in effect even when users are off the corporate network, such as when working from home or traveling. Because filtering occurs locally on the device, agent-based methods often provide more granular control, such as filtering HTTPS traffic or applying different filtering rules for different applications. Agent-based filtering can also provide detailed reporting and analytics.

Centralized Web Filtering

In a centralized approach, a proxy server performs content filtering by analyzing web requests from users and determine whether to permit or deny access based on established policies. A centralized proxy server acts as an intermediary between end users and the internet. Beyond blocking unwanted or harmful content, a centralized proxy server can also perform detailed logging and reporting of web activity.
A centralized proxy server employs various techniques to protect web traffic and ensure the safety of an organization's network. Examples of these techniques include:

- URL scanning—involves blocking URLs known to host malicious content, are inappropriate, or violate the company's policies.
- Content categorization—websites can be classified into various categories like social networking, gambling, webmail, and many others.
- Block rules—organizations can define rules to allow or deny access based on things like content categories, downloads of file types like .exe, or certain domains.
- Reputation-based filtering—leverages continually updated databases that score websites based on their observed behavior and history. Sites known for hosting malware, engaging in phishing attacks, or distributing spam, for instance, would have a poor reputation score and could be automatically blocked.

Web Filtering Challenges

Content and web filtering are not without potential issues and challenges. Three common challenges include:

- Overblocking which occurs when the filter is too restrictive, inadvertently blocking access to legitimate and useful websites and negatively impacting employee productivity.
- Underblocking which occurs when the filter allows access to potentially harmful or inappropriate websites.
- Handling of encrypted traffic like HTTPS. Without proper configuration, web filters may be unable to inspect encrypted traffic, representing most modern web traffic.

DNS Filtering

Domain Name System (DNS) filtering is a technique that blocks or allows access to specific websites by controlling the resolution of domain names into IP addresses. It operates on the principle that for a device to access a website, it must first resolve its domain name into its associated IP address, a process managed by DNS. When a request is made to resolve a website URI, the DNS filter checks the request against a database of domain names. If the domain is associated with malicious activities or is on an unapproved list for any reason, the filter blocks the request, preventing access to the potentially harmful website.

DNS filtering is highly effective for many reasons. A few are listed below:

- It provides a proactive defense mechanism, blocking access to known phishing sites, malware distribution sites, and other malicious online destinations.
- It can help enforce an organization's acceptable use policies (AUPs) by blocking access to inappropriate or distracting websites and ensuring that the internet is used responsibly and productively.
- It can protect all devices connected to a network, including IoT devices, providing an extra layer of security.
- It is a simple solution that is easy to implement and presents minimal risk, making it a cost-effective security control suitable for networks of any size.

While DNS filtering is highly effective, it must be combined with other security measures for comprehensive protection.

5.2.4 Security Zones (Lesson Video)

Transcript:

In this lesson, we're going to discuss security zones. Security zones are portions of the network that have specific security concerns or requirements. All devices within the same zone have the same security access and protection needs.

These zones are often separated by a traffic control device, such as a firewall or a router, that filters incoming and outbound traffic.
For example, you can define a zone that includes all hosts on your private network that need to be protected from the internet. You can also define a zone within your network for controlled access to specific servers that hold sensitive information.
There are several types of security zones that can be implemented.
Let's start with a security zone known as a screened subnet. A screened subnet is used as a public-facing accessible network. It acts as a buffer and a barrier to your internal production network, which is accessible from the internet. A screened subnet can be configured using either one or two firewalls and acts as the middleman between the internet and your internal network.
For example, a screened subnet is often used to allow employees access to critical network components, such as email servers, VPN connections, and internal web servers, from outside the company. Screened subnets also provide customer access to a company's products, such as software.
To do this, the firewall allows traffic from the internet into the screened subnet and enables those resources to access resources on the production network. The servers placed in the screened subnet are typically configured as a bastion host. A bastion host is a server that's specifically designed and configured to withstand attacks. In general, they're specific-purpose systems that host a single application and are hardened to resist attacks.
Another type of security zone is a wireless network zone. A wireless network zone is a broadcasted network connection used within an organization where users don't need a physical connection to a network port to connect to internal resources.
Instead, they use a wireless connection on their device to connect to a wireless access point. For internal users, the effect is the same as if the user's device has a physical connection and provides access to internal resources.
Several organizations also create a guest network that allows non-internal staff to access the internet. In some cases, guests may also be granted access to a few limited internal resources.

A honeynet contains servers designed to trap attackers. The servers, or honeypots, contain intentional vulnerabilities to lure attackers. Attackers believe they're accessing vulnerable systems, but they're really inside a trap. Honeypots are designed to mimic production servers that would be found in a typical organization. Of course, they contain no real company information, so there's little risk of exposure. The honeypots track an attacker's activity and can generate extremely useful security information.

An ad hoc network is a decentralized network that allows connections without a traditional base station or router. It allows users to connect two or more devices directly to each other for a specific purpose. Ad hoc is a Latin term that translates to "for this," and ad hoc networks are usually temporary solutions for a specific task.

Ad hoc connections can be wired or wireless. If you're connecting two systems with an Ethernet cable, it must be a crossover cable. An ad hoc network can also be created using Wi-Fi. Of course, the Wi-Fi parameters and protocols must match and be compatible with each other, including the SSID, security, and encryption settings.

Network Address Translation, or NAT, provides global IP address conversion for the entire network and is usually performed by your internet gateway, or NAT server. NAT allows network administrators to manage IP addresses within the internal network using whatever private IPs they want. When a user needs to access the internet, the NAT server then translates these internal private IP addresses to public IP addresses that the internet can recognize.

The organization's internet service provider can supply internet routable IP addresses and assign them to the organization. When information returns, the internet gateway converts the public address back to the private address and sends it on to the private network and user.

That's it for this lesson. In this lesson, we discussed several types of security zones. First, we discussed screened subnets, where servers are placed for public access. Next, we discussed wireless zones, which allow employees to access internal resources using Wi-Fi and provide guest networks that limit access. Then, we talked about honeynets, which are used to trap attackers and gather information. We ended this lesson by discussing ad hoc networks, which are used for temporary connectivity between two or more devices, and NAT servers, which convert private IP addresses to public IP addresses.

5.2.5 Security Zone Facts

This lesson covers the following topics:

- Security zones
- Security zone networks
- Common security zones

Security Zones

Security zones are portions of the network or system that have specific security concerns or requirements. All devices with the same zone have the same security access and security protection needs. These zones are often separated by a traffic control device, such as a firewall or a router, that filters incoming and outbound traffic. For example, you can define a zone that includes all hosts on your private network protected from the internet. You can also define a zone within your network for controlled access to specific servers that hold sensitive information.

Security Zone Networks

The following table lists the types of networks found in your security zones:

Network Type	Description
Wireless	A wirelessly broadcasted network is used on most internal networks so that internal users do not require a physical connection to a router or switch.

Guest	A guest network at an organization often grants internet access only to guest users, but it also has some type of firewall to regulate that access. There could be limited internal resources made available on a guest network. Normally, it is just a way for guests to access the internet without being allowed on the intranet or internal network.
Honeynet	A honeynet is a special network created to trap potential attackers. Honeynets have vulnerabilities that lure attacks so that you can track their actions and protect your real network. Honeynets can generate extremely useful security information.
Ad hoc	An ad hoc network is a decentralized network that allows connections without a traditional base station or router. It allows users to connect two or more devices directly to each other for a specific purpose.

Common Security Zones

The following table lists common zones:

Zone	Description
Screened subnets	A screened subnet is used as a public-facing accessible network. It acts as a buffer and a barrier to your internal production network, which is accessible from the internet. A screened subnet can be configured using either one or two firewalls and acts as the middleman between the internet and your internal network. For example, a screened subnet is often used to allow employees access to critical network components, such as email servers, VPN connections, and internal web servers, from outside the company. Screened subnets are also used to provide customer access to a company's products, such as software. To do this, the firewall allows traffic from the internet into the screened subnet and enables those resources to access resources on the production network. The servers placed in the screened subnet are typically configured as a bastion host. A bastion host is a server that's specifically designed and configured to withstand attacks. In general, they are specific-purpose systems that host a single application and are hardened to resist attacks.
Intranet	An intranet is a private network (LAN) that employs internet information services for internal use only. For example, your company network might include web servers and email servers that are used by company employees.
Extranet	An extranet is a privately controlled network distinct from the intranet but located between the internet and a private LAN. An extranet is often used to grant resource access to business partners, suppliers, and even customers outside of the organization.
Wireless	A wireless zone is a broadcasted network connection used within an organization. Users do not need a physical connection to a network port to connect to the intranet or internal resources. Instead, they use a wireless connection on their device to connect to a wireless access point.

5.2.6 Configure Network Security Appliance Access (Demo Video)

Transcript:

In this demonstration, we're going to secureÂ access to our LAN and WAN interfaces and configure IP addresses to them. We're also going to work with user accounts on a network security appliance.Â We're using a pfSense securityÂ appliance, which has a graphical interface forÂ configuration that's accessible through a web browser.

I'm here at the login screen, and I'll sign in with the default username of â€˜admin'. For this pfSense device, the default password is â€˜pfsense'. I'll click Sign In to log on.
The first thing I want to do is configure the IP address on my WAN and my LAN. We'll start with the LAN settings. Go up to Interfaces > LAN, and that takes us to the LAN configuration page. If I look here, the IPv4 Configuration Type is set to Static IPv4. DHCP is the other option there. I need to change it to a different IP address. I want to change it to â€˜10.10.10.1', so I'll do that now. Now I want to confirm that my subnet is set to the /24 subnet mask. That's all I need to do on this page. I'll scroll down and click Save.
The page refreshes and says that the LAN configuration has changed, and I need to click Apply Changes for it to take effect. It reminds me to change my DHCP server configuration if needed. Since my IP address is in the same subnet range as it was before, my DHCP configuration is still okay. Just make sure the IP you assigned is not in the scope of IPs you've configured to be leased. Click Apply Changes.
Now, if you look up at my address, you'll see that it says 10.10.10.254. When I log back in, I need to go to the new address that I configured. I've found that you sometimes get a token error if you just try to log back in, so I'm going to close my web browser and open it back up.
Now let's type in '10.10.10.1' to get back to our Sign In page. I'll put my credentials back in here and click Sign In. I'll go back to my LAN interface, and you can see that my changes are there.
Now let's configure our WAN settings. Once again, I'll go to Interfaces. But this time, I'll select WAN from the list. Right now, I'm getting an IP for my WAN interface via DHCP. I want to configure this with a static IP, so I'll change that to Static IPv4 from our IPv4 configuration type.
Now, for an IP address, I'll give it â€˜192.168.25.254'. Right away, you might be saying, "Hey, that's a Class C Private IP address." If so, you're absolutely correct. I'm on a test network that's connected to my regular network, which is using the 192.168 scheme. I need to set my subnet mask to /24 here. For this demo, I'm not going to configure DHCP version 6. I'm going to leave everything else set to the defaults and click Save to continue.
I get a few reminders, just like we did with the LAN configuration. I'll click Apply Changes.
Now I want to configure the DNS that my WAN will use. A lot of devices will have those settings right here, with the WAN settings. But with this device, we need to go up to System > General Setup.
I'm going to add a few DNS servers here. But first, let's take a look at the current status of our WAN interface. To do that, I'll go to Status > Interfaces. Here, we see that my DNS server is using the home address of 127.0.0.1. IPv4 network standards reserve the entire address block 127.0.0.0/8 for loopback purposes, so this is normal. I want to change those to a different DNS, so let's go back to System > General Setup. Beneath DNS Server Settings, I'll put in the Google DNS IP of â€˜8.8.8.8'. Now I'll click the Add DNS Server button, and this time, I'll put in the other Google DNS IP of â€˜8.8.4.4'.
I'm going to check this box, Disable DNS Forwarder. For this demo, I don't want my local host, 127.0.0.1, to be used as my first DNS server. Checking the box will remove the local host DNS. I'm only doing this to demonstrate what happens when you check the box. I'll scroll down to save these settings.
Now let's look at the status of our interfaces. We can see our DNS settings are now set to the Google DNS servers. That wraps up configuring IP interfaces.
Now let's add an additional user to our pfSense appliance. To add a user, we need to go up to System > User Manager. You can see here that we have just one user, the admin user. It's good practice not to use the admin account for normal maintenance. You usually want to create a separate account for each person that might mange the device, collect logs, and so on. We can also track the logon and logoff events of the users in case something happens.
I'll add the user by going over and clicking the Add User button. I'll enter in a username of â€˜Rachel McGaffey'. I'll give Rachel a password, and I'll put the password in a second time to confirm it. Now, under Full Name, I'll enter in â€˜Rachel McGaffey'. That's really all that's required. But if I wanted to, I could check this box to customize the GUI for Rachel's account. Let's take a look at some of those options.
After I check the box, I get some more options. I can change the theme for her account from this dropdown list. We can change the top navigation of the page here. We can change how the hostname appears in the menu and change the

amount of dashboard columns. If we want things like WAN and LAN to be sorted alphabetically, we can change that too. If we have groups set up, we can change the memberships here. But we won't make any changes to these right now. Let's come down and click Save.
One thing we don't want to happen with our security appliance is have someone logged in to it make configuration changes. If a user gets called away from his or her desk and doesn't sign out, anyone can sit down at their computer and start messing around. To mitigate this risk, we want to configure Session Timeout. I'll click on the Settings tab. Under Session Timeout, I'm going to change this to 10 minutes. That is actually a long time, and you might want to set this to a lower amount of time. Having no timeout is a huge security risk. You'll have to judge your circumstances and decide what session timeout settings are best.
The last thing I want to do in this demo is change the default admin password. Default passwords for devices are easy to find with a simple web search. It's such a problem, some US states and other governments have forced many tech manufacturers to make the devices they produce have random passwords. Even if your device comes with a random password already configured, those passwords are often printed on the device somewhere, so it's always recommended that you change itâ€"that should be the very first thing you do.
To change this password for the admin account, I'll click on the edit icon here. I'll type in my new password. Now I'll type it a second time to confirm it. Come down and click Save, and that's all there is to it.
That's it for this demonstration.Â In this demo, we configured a network security appliance.Â We configured the LAN and WAN IP addresses. We changed the WAN DNS. We created a second account on the device. We configured a session timeout. We ended the demo by changing the defaultÂ admin password to something more secure.

5.2.7 Configure a Security Appliance (Simulation)

Scenario

You are an IT security administrator for a small corporate network. To increase security for the corporate network, you have installed the pfSense network security appliance in your network. Now you need to configure the device.

In this lab, your task is to configure pfSense as follows:

- Sign in to pfSense using the following case-sensitive information:
  - URL: **198.28.56.22**
  - Username: **admin**
  - Password: **P@ssw0rd**
- Configure the DNS servers as follows:
  - Primary DNS server: **163.128.78.93** - Hostname: **DNS1**
  - Secondary DNS server: **163.128.80.93** - Hostname: **DNS2**
- Configure the WAN IPv4 information as follows:
  - Enable the interface.
  - Use a static IPv4 address of **65.86.24.136/8**
  - Add a new gateway using the following information:
    - Type: **Default gateway**
    - Name: **WANGateway**
    - IP address: **65.86.1.1**

Explanation

Complete this lab as follows:

1. Access the pfSense management console.
   a. From the taskbar, select **Google Chrome** .
   b. Maximize the window for better viewing.
   c. In the address bar, type **198.28.56.22** and then press **Enter** .
   d. Sign in using the following case-sensitive information:
      - Username: **admin**

        - Password: **P@ssw0rd**
    e. Select **SIGN IN** or press **Enter** .
2. Configure the DNS Servers.
    a. From the pfSense menu bar, select **System** > **General Setup** .
    b. Under *DNS Server Settings* , configure the primary DNS Server as follows:
        - Address: **163.128.78.93**
        - Hostname: **DNS1**
        - Gateway: **None**
    c. Select **Add DNS Server** to add a secondary DNS Server and then configure it as follows:
        - Address: **163.128.80.93**
        - Hostname: **DNS2**
        - Gateway: **None**
    d. Scroll to the bottom and select **Save** .
3. Configure the WAN settings.
    a. From the pfSense menu bar, select **Interfaces** > **WAN** .
    b. Under General Configuration, ensure **Enable interface** is selected.
    c. Use the *IPv4 Configuration Type* drop-down to select **Static IPv4** .
    d. Under Static IPv4 Configuration, in the IPv4 Address field, enter **65.86.24.136** .
    e. Use the IPv4 Address subnet drop-down to select **8** .
    f. Under Static IPv4 Configuration, select **Add a new gateway** .
    g. Configure the gateway settings as follows:
        - Default: Select **Default gateway**
        - Gateway name: Enter **WANGateway**
        - Gateway IPv4: **65.86.1.1**
    h. Select **Add** .
    i. Scroll to the bottom and select **Save** .
    j. Select **Apply Changes** .

5.2.8 Configure Network Security Appliance Access (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. You need to secure access to your pfSense appliance, which is still configured with the default user settings.

In this lab, your task is to:

- Change the password for the default pfSense account from P@ssw0rd to **1w0rm4b8.**
- Create a new administrative user with the following parameters:
    - Username: **zolsen**
    - Password: **St@yout!**
    - Full Name: **Zoey Olsen**
    - Group Membership: **admins**
- Set a session timeout of **15** minutes for pfSense.
- Disable the webConfigurator anti-lockout rule for HTTP.

Access the pfSense management console through Google Chrome using: **http://198.28.56.22**

- Default username: **admin**
- Password: **P@ssw0rd**

Explanation

Complete this lab as follows:

1. Access the pfSense management console.
    a. From the taskbar, select **Google Chrome** .
    b. Maximize the window for better viewing.
    c. In the Google Chrome address bar, enter **198.28.56.22** and then press **Enter** .
    d. Enter the pfSense sign-in information as follows:
        - Username: **admin**
        - Password: **P@ssw0rd**
    e. Select **SIGN IN** .
2. Change the password for the default (admin) account.
    a. From the pfSense menu bar, select **System > User Manager** .
    b. For the admin account, under Actions, select the ***Edit user*** icon (pencil).
    c. For the Password field, change to **1w0rm4b8** .
    d. For the Confirm Password field, enter **1w0rm4b8** .
    e. Scroll to the bottom and select **Save** .
3. Create and configure a new pfSense user.
    a. Select **Add** .
    b. For Username, enter **zolsen** .
    c. For the Password field, enter **St@yout!** .
    d. For the Confirm Password field, enter **St@yout!**
    e. For Full Name, enter **Zoey Olsen** .
    f. For Group Membership, select **admins** and then select **Move to Member of list** .
    g. Scroll to the bottom and select **Save** .
4. Set a session timeout for pfSense.
    a. Under the *System* breadcrumb, select **Settings** .
    b. For Session timeout, enter **15** .
    c. Select **Save** .
5. Disable the webConfigurator anti-lockout rule for HTTP.
    a. From the pfSense menu bar, select **System** > **Advanced** .
    b. Under webConfigurator, for Protocol, select **HTTP** .
    c. Select **Anti-lockout** to disable the webConfigurator anti-lockout rule.
    d. Scroll to the bottom and select **Save** .

5.2.9 Deceptive and Disruption Technology (Lesson Video)

Transcript:

In this lesson we'll discuss attack deceptions. By now you should have a clear understanding of what an attack is and what it looks like. Threat agents are working hard to find weaknesses in your defenses. They will exploit any vulnerability to gather information from an organization's network system.

Knowing this, organizations can set up data or add devices to their networks that lures unsuspecting attackers into attacking devices. The attacking devices are designed to trap attackers and prevent them from exfiltrating any useful information.
The first item we'll look at is the honeypot. As the name implies, if you look inside, you'll find will find something sweet. A honeypot is a network set up to fool an attacker into thinking he's attacking a network device critical to your company; however, it's simply a ruse. The attacker can attack the honeypot as long as he wants and get as deep as he wants; however, there is no prize or honey in the end.
A honeypot can be a host, a service on a host, a network device, a virtual entity, or even a single file set up to attract attackers to a secure area away from an organization's real network. Even better, while it's distracting the attacker, you can monitor the malicious activity to learn what the attacker is trying to do.
For example, if an attacker attempts an exploit or upload a rootkit or Trojan to a server, the honeypot environment will safely store those files for malware collection and analysis.

There are several types of honeypots that are designed to mimic and trap different types of attacks. Each honeypot falls into one of two categories: research or production. Research honeypots are typically used by research, military, or government organizations to gather information about the motives and tactics of the attackers. They are also used for educational purposes. Production honeypots are used, primarily by private and public organizations.

Let's look at a few examples of production honeypots.

First, you have the SSH honeypots. SSH honeypots are designed to trap brute force attacks that attempt to guess the password of a targeted device. Perhaps more importantly, this type of honeypot captures the complete shell history performed by the attacker. Next, we have HTTP honeypots. This type of honeypot is used to trap different types of web vulnerabilities such as SQL injections, file insertions, and web application attacks.

We also have WordPress honeypots which are designed to attract spam and adware attacks. An email honeypot is used to attract SMTP-based attacks such as open relay, DDoS, and spam attacks. And lastly, there are several types of honeypots designed for TCP/IP-based protocol attacks.

To make things more enticing to the attackers, many companies will create a decoy network and fill it with one or more honeypots. This is known as a honeynet. The information gained from a honeynet attack serves to strengthen an organization's defense against future attack.

Similar in concept to a honeypot, a honeyfile is a single file setup to entice and trap attackers and to figure out what they're trying to do. These types of files are typically set up on a file server, and when accessed, an alarm is set off notifying that administrators of the attack. These types of alarms are often sent by email directly to the user, which can also be easily seen on a phone and almost instantly interpreted. Honeyfiles are named so the hacker will try and open or execute them. For example, files with names such as passwords.txt or IP Assignments.xls are prized targets for hackers to gain useful information about a user or system.

When it comes to honeyfiles, attackers are typically interested in four types of files. First, are the informational files. These are the types of files that contain information about access other systems, like a password.txt file, and security files like Microsoft word file named vpn_instructions. The second type are application or program files that an attacker may run, but the normal user wouldn't, such as a complier. Third, hackers love to attack files that contain evidence of an attack such as log files. The last type of file commonly targeted by attackers are files that might contain intellectual property such as a forecast, credit card numbers, and product roadmaps.

As you can see, honeyfiles protect a wide variety of files and systems including email and email attachments. Although we've pointed out the four major types of file that attack hackers, keep in mind that any file can be used as a honeyfile. Honeyfiles also work with network intrusion detection systems, or NIDS, and help to prevent false positives.

The other nice thing about honeyfiles is that if end users are properly trained, they can create their own honeyfiles and set up their own alarms and alerts. Enabling users makes the detection process easy and deployment effectiveness is high.

Since users create their own honeyfiles, they receive the alarm if they accidentally access one of their own monitored files. This virtually eliminates false positives that have be investigated. This saves time, effort, and energy for system management charged with keeping the network secure. Further, this process supports a defense-in-depth strategy.

Now let's shift gears and discuss DNS sinkholes.

Simply put, a DNS sinkhole is used to provide false information for DNS queries. Although this may sound counterintuitive, providing false information can help to prevent a denial-of-service or distributed denial-of-service attack. DNS sinkholes are an effective measure that blocks malicious traffic and blocks bots from successfully taking a site down using a DDoS attack.

It's important to understand that a DNS sinkhole is not a product. Instead, it's a method used by system management to fool an attacker by providing bogus information. For example, suppose a bot attempts to use DNS to forward information. The DNS server returns fake information and the bot can't continue its malicious activity. As another example, supposed an attacker sends a malicious email to an unsuspecting user. Should the user click the link in the email, the DNS sinkhole would prevent the attack by resolving the DNS query to a false site.

Creating a DNS sinkhole is not without its limitations. To work correctly, the malware must use the organization's DNS server. If it uses a public DNS server or its own DNS server, this won't work. You can mitigate this by configuring firewalls to block DNS queries going outside the perimeter.

It's also important to understand that a sinkhole doesn't prevent malware and can't prevent malware execution or remove malware from the system. However, the DNS server will be configured with potential malware sites.

Creating a sinkhole can take time. Even if obtained from trusted internet sources, its possible legitimate sites are used to forward malicious content. This could result in restrictions to legitimate websites. Therefore, sinkholes should be internal only. If attackers can access the DNS server externally, they can change entries in the sinkhole and use it to their benefit.

While many threat agents are looking for open and available vulnerabilities, others are looking to infiltrate a network's defenses and either copy intellectual property or damage internal systems. It's truly a cat and mouse game played by

attackers and network managers. Attackers are always trying new techniques to perform their reconnaissance and discover weakness they can exploit. One method used by system managers is the concept of fake telemetry. Telemetry is defined as the collection of data at remote points and their automatic transmission to receiving equipment for monitoring. As such, organizations provide false data, meant to deceive the attacker. This type of fake information might include fake credentials or fake IP address information. It might be an internal report detailing vulnerable systems that require updates when, they are really honeypots designed to capture attacker information.

The primary reason for providing fake information is to get the attacker to act on the false information. Since the information is controlled by system managers, they can utilize alerts, alarms, and logs to track the techniques and tools used by an attacker. The tracked information allows system management to tune perimeter equipment such as firewalls, NIDS, NIPS and other protections.

That's it for this lesson. In this video, we discussed deception techniques used by companies to protect themselves from attackers. We discussed such techniques as using honeypots, honeyfiles, and honeynets to lure attackers into fake areas. We also discussed the benefits of using DNS sinkholes. And we ended this lesson outlining the benefits of companies using fake telemetry.

5.2.10 Deceptive and Disruption Technology Facts

Deception and disruption technologies are cybersecurity resilience tools and techniques used to increase the cost of attack planning for the threat actor.

This lesson covers the following topics:

- Honeypots
- Honeynets and honeyfiles
- DNS sinkholes
- Fake telemetry

Honeypots

Honeypots are decoy systems that mimic real systems and applications. They are designed to allow security teams to monitor attacker activity and gather information about their tactics and tools. A honeypot can be a host, a service on a host, a network device, a virtual entity, or even a single file set up to attract attackers to a secure area away from an organization's real network. Even better, while it is distracting the attacker, you can monitor the malicious activity to learn what the attacker is trying to do.

There are several types of honeypots that are designed to mimic and trap different types of attacks. Each honeypot falls into one of two categories: research or production. Research honeypots are typically used by research, military, or government organizations to gather information about the motives and tactics of the attackers. They are also used for educational purposes. Production honeypots are used primarily by private and public organizations. Examples include the following:

- SSH honeypots — SSH honeypots are designed to trap brute force attacks that attempt to guess the password of a targeted device. Perhaps more importantly, this type of honeypot captures the complete shell history performed by the attacker.
- HTTP honeypots — HTTP honeypots are used to trap different types of web vulnerabilities, such as SQL injections, file insertions, and web application attacks.
- WordPress honeypots — WordPress honeypots are designed to attract spam and adware attacks.
- Email honeypots — Email honeypots are used to attract SMTP-based attacks such as open relay, DDoS, and spam attacks.
- Various other honeypots designed for TCP/IP-based protocol attacks.

Honeynets and Honeyfiles

To make things more enticing to the attackers, many companies will create a decoy network and fill it with one or more honeypots. This is known as a honeynet. The information gained from a honeynet attack serves to strengthen an organization's defense against future attacks.

A honeyfile is a single file setup to entice and trap attackers and to figure out what they are trying to do. These types of files are typically set up on a file server, and when accessed, an alarm is set off, notifying administrators of the attack. These types of alarms are often sent by email directly to the user, which can also be easily seen on a phone and almost instantly interpreted.

Honeyfiles are named so the hacker will try and open or execute them. For example, files with names such as passwords.txt or IP Assignments.xls are prized targets for hackers to gain useful information about a user or system. Four of the most common types of files used to attract an attacker include the following:

- Information files — These are the types of files that contain information about accessing other systems, like a password.txt file, and security files, like a Microsoft Word file named vpn_instructions.
- Application or program files — These are files that a normal user would not run but an attacker might, such as a compiler.
- Log files — An attacker may be attracted to log files, thinking they might contain evidence of an attack or other helpful information.
- Intellectual property files containing information such as forecasts, credit card numbers, and product roadmaps.

Keep in mind the following facts about honeyfiles:

- Any file can be used as a honeyfile.
- Honeyfiles work with network intrusion detection systems (NIDs) and can help prevent false positives.
- If end users are properly trained, they can create their own honeyfiles with their own alarms and alerts.
    - This makes the detection process easy and deployment effectiveness high.
    - Since users create their own honeyfiles, they receive the alarm if they accidentally access one of their own monitored files.
        - This virtually eliminates false positives that have to be investigated.
        - This saves time, effort, and energy for system management charged with keeping the network secure.
        - This process supports a defense-in-depth strategy.

Honeytokens are false credentials, login credentials, or other data types used to distract attackers, trigger alerts, and provide insight into attacker activity.

By deploying honeypots, honeynets, honeyfiles, or honeytokens, organizations can detect and monitor attacks, gather intelligence about attackers and their methods, and proactively defend against future attacks. These tools can also provide an additional layer of defense by diverting attackers' attention away from real systems and applications, reducing the risk of successful attacks.

DNS Sinkhole

A DNS sinkhole is a temporary DNS record that redirects malicious traffic to a controlled IP address. Simply put, a DNS sinkhole is used to provide false information for DNS queries. Providing false information can help to prevent a denial-of-service or distributed denial-of-service attack. DNS sinkholes are also an effective measure that blocks malicious traffic and blocks bots from successfully taking a site down using a DDoS attack.

Creating a DNS sinkhole is not without its limitations. To work correctly, the malware must use the organization's DNS server. If it uses a public DNS server or its own DNS server, this will not work. You can mitigate this by configuring firewalls to block DNS queries going outside the perimeter. A sinkhole does not prevent malware and cannot prevent malware execution or remove malware from the system.

Creating a sinkhole can take time. Even if obtained from trusted internet sources, it's possible legitimate sites are used to forward malicious content. This could result in restrictions to legitimate websites. Therefore, sinkholes should be internal only. If attackers can access the DNS server externally, they can change entries in the sinkhole and use it to their benefit.

Fake Telemetry

Fake telemetry is a deception strategy that returns spoofed data in response to network probes. Telemetry is defined as the collection of data at remote points and their automatic transmission to receiving equipment for monitoring. As such, organizations provide false data meant to deceive the attacker. This type of fake information might include fake credentials or fake IP address information. It might be an internal report detailing vulnerable systems that require updates when they are really honeypots designed to capture attacker information.

The primary reason for providing fake information is to get the attacker to act on the false information. Since the information is controlled by system managers, they can utilize alerts, alarms, and logs to track the techniques and tools used by an attacker. The tracked information allows system management to tune perimeter equipment such as firewalls, NIDS, NIPS, and other protections.

5.2.11 Detect Malicious Network Traffic with a Honeypot (Demo Video)

Transcript:

Reconnaissance tools like nmap can find vulnerable systems on the internet. These tools scan ranges of IP addresses for systems with interesting ports open, such as Telnet, Remote Desktop, SSH, and FTP. Many ethical hackers like to know who's gathering this kind of information, and there's a simple way to record reconnaissance attempts. In order to record these network scans, we need a computer whose sole purpose is to listen for connection attempts on interesting ports, then log the data about each attempt. This kind of system is called a honeypot. It looks appealing and hackable from the outside, but it's actually recording data about every remote user that attempts to connect. Today, we'll look at how to set up an extremely basic honeypot using a tool called Pentbox.

Pentbox is available online, but many people upload copies of the original program. When downloading tools, especially hacking tools, you have to be certain of the tool's integrity so you don't download malicious software. I've already downloaded the tool and put it in a folder named pentbox, so let's open that now. From my Kali Linux machine, I'll open a terminal and navigate to 'cd pentbox/'. I'll list the folder's contents to make sure everything is there.
Pentbox is a Ruby script. Since Ruby is already installed on Kali, we just need to run the script. To do that, I'll type './pentbox.rb' and press Enter.
When it opens, we can see a variety of options. Pentbox does a lot of things. But right now, we're only interested in the honeypot utility, which is located under Network Tools. First, I'll select number 2, and then we'll select Honeypot, number 3.
After we've selected the Honeypot utility, we're presented with the option to use either fast or manual configuration. I'm going to select number 1, the Fast Autoconfiguration option, since the manual configuration only lets us change which port the tool listens on. It also lets us set the message to return to the requesting machine.
As soon as we press Enter, we're told that the honeypot has started running on port 80. Let's try connecting to it with Firefox to see what happens. I'll open Firefox and type 'localhost' in the address bar. We can see that the honeypot returns back a web page that tells us access has been denied. Interestingly, the autoconfiguration option returns a line with a date and time. This doesn't change; actually, it represents the time the tool was started. Now that we've connected to the page, let's see what Pentbox says.
Okay, I'm back on my Kali Linux system. It looks like the honeypot utility is telling us quite a bit of information. The first line of each entry shows that there was a connection attempt and tells us what the IP address of the connecting machine was, as well as the connection time. After that, the honeypot utility tells us the header information that was received when the connection was made. This information includes both the web browser that was used and information about the type of operating system that connected.
So, there's a really simple overview of a honeypot's function. This is just a taste of what you can accomplish- there are many more robust tools available that can log information about connections on many different ports and protocols simultaneously, which is a goldmine of data that will help you keep your network safe.

5.2.12 Configure Load Balancer (Demo Video)

Transcript:

A load balancer can be either a software or a hardware appliance that distributes the load across multiple servers using TCP or UDP. Some examples could include websites, web services, email, print servers, and much more. Not only can it serve the purpose of load balancing, but it can also be used for failover. Today, we're going to be using a software load balancer called Kemp. This is an easy-to-install software load balancer deployed to VMware with an .ovf template.

One of the first things you must do is configure a Virtual Service. The virtual service will have an IP address associated with it. This is simply the IP address that users will connect to instead of the real server addresses where load balancing occurs. Our Virtual Address will be 192.168.30.31, which is outside our DHCP zone. Port 80 will be used since this is just an HTTP website. We can also add a Service Name so we know what this is for. Templates are useful because they help predefine some settings related to the type of service you're load balancing. These can be obtained from Kemp's website, as you can see here. This example shows Apache and Tomcat templates, but as we scroll up, there are many different types of services that we can load balance with Kemp. Let's go back to our load balancer. We'll select the Apache HTTP template and then click Add this Virtual Service. When this service is added, there are several options available, such as quality of service or limiting. These options can help ensure that the service you're load balancing isn't overloaded. The default is 0, which means off. Advanced settings can also assist with tasks such as redirecting or enabling caching.
At the bottom is where we'll add our real servers. These are the actual web servers that we'll be load balancing. When we click Add New, we'll input our first web server, 192.168.30.10, and then click Add This Real Server. Our second one will be 192.168.30.11. Both web servers will use port 80 since they're serving only HTTP websites. The Forwarding method will remain as NAT since that's the only option. The Weight can be adjusted if you want to send more traffic to one server over the other. For now, we've left them both at 1000. One useful feature is that when you hover over these options, it provides a description to help you configure them. Connection Limit and Connection Rate Limit can be set, but we won't be doing that today. Click Add This Real Server. After clicking OK, you can see that we have two real servers now configured for this virtual service.
We need to make some more changes for this virtual service to work correctly. If we go to View/Modify Services, we can see the virtual service we've created. Clicking Modify on the right will allow us to go back into the settings. Down at the bottom, we need to set the URL we're looking to load balance. Since our web servers don't have a website at the root directory, we need to add /test.html to the URL area. Next, click Set URL.
Now, we're ready to test out load balancing with our websites. If we open a new tab, we can go to http://192.168.30.10/test.html. This will show us our first web server, and if we modify the address and change the 10 to 11, we'll access our second web server. In a real-world scenario, the websites on both servers would be identical; however, we've made them different so you can tell which server you're on during load balancing. Now that we know our web servers work, let's use our virtual IP address. If we go back to our load balancer, it can detect that both web servers are online. Our virtual IP will be 192.168.30.31. Let's go there. When we type it in, you can see that it's directing us to webserver-svp01.
A common thing we can do with a virtual IP is associate it with a DNS name. On our Pi-hole DNS server, we can scroll down and go to Local DNS and DNS Records. You can already see that we have records for our web servers, so let's add a record for our virtual IP. The domain will be testwebsite.com, and our virtual IP address is 192.168.30.31. After clicking Add, it will add it to the list of records. Let's test this out. In our web tab, we'll remove the IP address and enter testwebsite.com, then click Enter. Great, it still shows webserver-svp01.
To make things more interesting, let's force traffic to go to webserver-svp02 since we have maintenance to attend to on webserver-svp01. If we go back to our load balancer tab and Modify our virtual service, we can go down to our real servers and click Disable on webserver-svp01. Once this is disabled, we can go back and test. Please note that many browsers cache websites, so in order for this change to be visible right away, we must open a new private browsing window. Now, if we enter our web address testwebsite.com/test.html, we can see that it's now showing webserver-svp02.
That's it for this demo. In this demo, we showed you how to configure a load balancer, add web servers to the load balancer, and test both IP and DNS with our web servers.

5.2.13 Practice Questions (Section Quiz)

q_sec_sol_appliance_secp8

You are the office manager of a small financial credit business. Your company handles personal financial information for clients seeking small loans over the internet. You are aware of your obligation to secure clients records, but the budget is an issue for your company.

Which item would provide the BEST security for this situation?

Answers:

- Network access control system
- Proxy server with access controls
- ***All-in-one security appliance**
- Firewall on your gateway server to the internet

Explanation:

An all-in-one security appliance would provide the best overall protection. All-in-one security appliances take up the least amount of space and require the least amount of technical assistance for setup and maintenance.

Security functions in an all-in-one security appliance can include the following:

- Spam filter
- URL filter
- Web content filter
- Malware inspection
- Intrusion detection system (IDS)

In addition to security functions, all-in-one security appliances can include the following:

- Network switch
- Router
- Firewall
- Tx uplink (integrated CSU/DSU)
- Bandwidth shaping

q_sec_sol_block_rules_secp8

A technician is deploying centralized web filtering techniques across the enterprise.

What stems from various factors such as the website's URL, domain, IP address, content category, or even specific keywords within the web content?

Answers:

- ***Block rules**
- Reputation-based filtering
- URL scanning
- Content categorization

Explanation:

Block rules stem from various factors such as the website's Uniform Resource Locators (URL), domain, Internet Protocol (IP) address, content category, or even specific keywords within the web content.

Reputation-based filtering leverages continually updated databases that score websites based on their observed behavior and history.

URL scanning is a method that protects web traffic as the proxy server examines the URLs requested by users and can block access to specific URLs known to host malicious content, are inappropriate, or violate the company's Internet usage policy.

Content categorization classifies websites into various categories, such as social networking, gambling, adult content, webmail, and many others.

q_sec_sol_content_01_secp8

You are implementing security at a local high school that is concerned with students accessing inappropriate material on the internet from the library's computers. The students use the computers to search the internet for research paper content. The school budget is limited.

Which content filtering option would you choose?

Answers:

- Block all content except for content you have identified as permissible.
- ***Restrict content based on content categories.**
- Block specific DNS domain names.
- Allow all content except for the content you have identified as restricted.

Explanation:

Restricting content based on categories would provide the most protection with the least amount of research and involvement.

All other options require research to identify specific content or websites, which could allow access to undesirable websites or prevent access to necessary websites.

q_sec_sol_content_02_secp8

A cyber group is reviewing its web filtering capabilities after a recent breach.

Which centralized web-filtering technique groups websites into categories such as social networking, gambling, and webmail?

Answers:

- ***Content categorization**
- Block rules
- Reputation-based filtering
- URL scanning

Explanation:

Content categorization classifies websites into categories such as social networking, gambling, adult content, webmail, and many others.

Block rules stem from various factors such as the website's Uniform Resource Locators (URL), domain, Internet Protocol (IP) address, content category, or even specific keywords within the web content.

Reputation-based filtering leverages continually updated databases that score websites based on their observed behavior and history.

URL scanning is a method that protects web traffic as the proxy server examines the URLs requested by users and can block access to specific URLs known to host malicious content, are inappropriate, or violate the company's internet usage policy.

q_sec_sol_fail_closed_open_secp8

Which of the following descriptions is true about fail-open and fail-closed configurations for security devices in the event of a failure?

Answers:

- Fail-open prioritizes confidentiality and integrity over availability, while fail-closed prioritizes availability over confidentiality and integrity.
- Fail-open means that access is blocked or that the system enters the most secure state available, while fail-closed means that network or host access is preserved, if possible.
- ***Fail-open means that network or host access is preserved, if possible, while fail-closed means that access is blocked or that the system enters the most secure state available.**
- Both fail-open and fail-closed prioritize confidentiality and integrity over availability.

Explanation:

In a fail-open configuration, the system maintains network or host access, if possible, in the event of a failure. In a fail-closed configuration, the system blocks access or enters the most secure state available in the event of a failure.

Fail-open prioritizes availability over confidentiality and integrity, while fail-closed prioritizes confidentiality and integrity over availability.

Fail-open means preservation of network or host access, if possible, while fail-closed means blocked access or the system entering the most secure state available.

q_sec_sol_fail_closed_secp8

You have configured a security device in your network to fail-closed.

Which of the following will happen when an attack occurs?

Answers:

- The device will preserve network or host access when it fails.
- ***The device will block access or enter the most secure state available when it fails.**
- The device will distribute network traffic across multiple servers when it fails.
- The device will act on behalf of a client when accessing resources over the internet when it fails.

Explanation:

In a fail-closed configuration, the system prioritizes confidentiality and integrity over availability. If a failure occurs, access is blocked or the system enters the most secure state available.

Preserving network or host access when the device fails describes a fail-open configuration, not fail-closed. In a fail-open scenario, the system prioritizes availability over confidentiality and integrity.

Distributing network traffic across multiple servers when the device fails describes the function of a load balancer, not a fail-closed configuration. A device configured to fail-closed will not distribute network traffic across multiple servers when it fails.

Acting on behalf of a client when accessing resources over the internet when the device fails describes the function of a proxy server, not a fail-closed configuration. A device configured to fail-closed will not act on behalf of a client when accessing resources over the internet when it fails.

q_sec_sol_load_balancer_secp8

What is the main role of a load balancer in network security?

Answers:

- ***To distribute network traffic across multiple servers.**
- To act as a gateway to access certain isolated systems.
- To capture packets for intrusion detection.
- To act on behalf of a client when accessing resources over the internet.

Explanation:

A load balancer distributes client requests across available server nodes in a farm or pool, optimizing performance and providing fault tolerance.

To act as a gateway to access certain isolated systems is the function of a jump server, not a load balancer.

To capture packets for intrusion detection is the purpose of a sensor, not a load balancer.

To act on behalf of a client when accessing resources over the internet is the function of a proxy server, not a load balancer.

q_sec_sol_proxy_01_secp8

A proxy server can be configured to do which of the following?

Answers:

- ***Restrict users on the inside of a network from getting out to the internet.**
- Allow all content except for the content you have identified as restricted.
- Block all content except for the content you have identified as permissible.
- Act as a unified threat security device or web security gateway.

Explanation:

Proxies can be configured to:

- Restrict users on the inside of a network from getting out to the internet.
- Restrict access by user or by specific website.
- Restrict users from using certain protocols.
- Use access controls to control inbound or outbound traffic.

- Shield or hide a private network to provide online anonymity and make it more difficult to track web surfing behavior.
- Cache heavily accessed web content to improve performance.

An internet content filter is software used to monitor and restrict content delivered across the web to an end user. Two types of configurations are commonly used, which are:

- Allow all content except for the content you have identified as restricted.
- Block all content except for the content you have identified as permissible.

All-in-one security appliances combine many security functions into a single device. All-in-one security appliances are also known as unified threat security devices or web security gateways.

q_sec_sol_web_filter_01_secp8

An organization needs to implement web filtering to bolster its security. The goal is to ensure consistent policy enforcement for both in-office and remote workers.

Which of the following web filtering methods BEST meets this requirement?

Answers:

- ***Deploying agent-based web filtering**
- Utilizing a centralized proxy server
- Implementing manual URL blocking
- Relying solely on reputation-based filtering

Explanation:

Agent-based web filtering involves installing a software agent on all devices. These agents communicate with a centralized server to obtain filtering policies and rules and apply them locally.

While a centralized proxy server effectively controls and monitors all inbound and outbound web content, it may struggle to enforce policies when employees cannot connect from the corporate network, such as when working remotely or traveling.

Manual URL blocking as a primary method is impractical, given the sheer volume of potentially harmful websites and the rapid development and creation of new websites.

Although reputation-based filtering is an effective technique and is part of a comprehensive web filtering strategy, solely relying on it would not provide all-around protection.

q_sec_sol_web_filter_02_secp8

An organization needs a solution for controlling and monitoring all inbound and outbound web content, analyzing web requests, blocking access based on various criteria, and offering detailed logging and reporting of web activity.

Which of the following solutions is the MOST suitable in this situation?

Answers:

- ***Centralized web filtering**
- Agent-based filtering

- Manual URL blocking
- Content categorization

Explanation:

Implemented typically through a proxy server, centralized web filtering controls and monitors all inbound and outbound web content. It can analyze web requests and block access based on Uniform Resource Locators (URLs), Internet Protocol (IP) addresses, content categories, or specific keywords.

Agent-based filtering provides detailed logging and can block access based on established policies but does not control and monitor all inbound and outbound web content at the network level.

Manual URL blocking isn't the most effective solution to control and monitor all inbound and outbound web content.

Content categorization as a standalone feature does not provide comprehensive control, monitoring, and logging of all web content.

q_sec_zone_ad_hoc_secp8

Which of the following BEST describes an ad hoc network?

Answers:

- A network that requires a physical connection to a router or switch.
- A network that is designed to trap potential attackers.
- ***A network that is decentralized and allows connections without a traditional base station or router.**
- A network that grants internet access only for guest users.

Explanation:

An ad hoc network is a decentralized network that allows devices to connect directly to each other without the need for a traditional base station or router. This type of network is often used for a specific purpose, such as a temporary connection for a meeting or event.

An ad hoc network does not require a physical connection to a router or switch. Instead, it allows devices to connect directly to each other.

A honeynet is designed to trap potential attackers, not to allow direct connections between devices.

A guest network typically provides internet access to visitors, but it does not allow direct connections between devices without a router or switch.

q_sec_zone_extranet_secp8

Which of the following is a privately controlled portion of a network that is accessible to some specific external entities?

Answers:

- ***Extranet**
- Intranet
- Internet
- MAN

Explanation:

An extranet is a privately controlled portion of a network that is accessible to some specific external entities. Often, those external entities are business partners, suppliers, distributors, vendors, or customers.

An intranet is a LAN that employs the technology of the internet (namely, TCP/IP, web servers, and email).

The internet is the global TCP/IP-based network that supports most web and email communications.

A metropolitan area network (MAN) is a LAN that is spread across several city blocks, across a business park, or across a campus.

q_sec_zone_honeynet_secp8

You want to create a collection of computers on your network that appear to have valuable data but actually store fake data that could entice a potential intruder. Once the intruder connects, you want to be able to observe and gather information about the attacker's methods.

Which feature should you implement?

Answers:

- ***Honeynet**
- NIDS
- NIPS
- Extranet

Explanation:

A honeypot is a device or virtual machine that entices intruders by displaying a vulnerable trait or flaw or by appearing to contain valuable data. A honeynet is a network of honeypots.

A network-based IDS (NIDS) is a dedicated device installed on a network that's used to analyze all traffic on the network. An NIPS is a network-based intrusion prevention system that can take actions in response to intrusion.

An extranet is a privately controlled network located between the internet and a private LAN, but distinct from both. An extranet is often used to grant resource access to business partners, suppliers, and even customers outside of the organization.

q_sec_zone_intranet_secp8

A company wants to set up a private network that employs internet information services for internal use only, including web servers and email servers that are used by company employees.

What type of network is the company planning to set up?

Answers:

- Extranet
- Internet
- ***Intranet**
- Honeynet

Explanation:

An intranet is a private network (LAN) that employs internet information services for internal use only. For example, a company network might include web servers and email servers that are used by company employees. This matches the scenario described in the question.

An extranet is a privately controlled network distinct from the intranet but located between the internet and a private LAN. An extranet is often used to grant resource access to business partners, suppliers, and even customers outside of the organization. This does not match the scenario where the company wants a network for internal use only.

The internet is a public network that includes all publicly available web servers, FTP servers, and other services. The internet is public because access is largely open to everyone. This does not match the scenario where the company wants a private network for internal use only.

A honeynet is a special network created to trap potential attackers. Honeynets have vulnerabilities that lure attacks so that you can track their actions and protect your real network. This does not match the scenario where the company wants a network for internal use only.

q_sec_zone_screened_subnet_secp8

You are a network security engineer for a large corporation. The company is planning to launch a new software product and wants to provide customer access to this product over the internet.

The company also wants to ensure that the internal production network remains secure.

Which type of common security zone would be the MOST appropriate to implement in this scenario?

Answers:

- Intranet
- Extranet
- ***Screened Subnet**
- Wireless

Explanation:

Screened subnet is the correct answer. A screened subnet is used as a public-facing accessible network. It acts as a buffer and a barrier to the internal production network, which is accessible from the internet. It can be configured using either one or two firewalls and acts as the middleman between the internet and the internal network. This makes it the most appropriate choice for this scenario.

Intranet is incorrect because an intranet is a private network that is used for internal purposes only. It would not be appropriate for providing customer access to a product over the internet.

Extranet is incorrect because an extranet is used to grant resource access to business partners, suppliers, and even customers outside of the organization. However, it is not specifically designed to act as a buffer between the internet and the internal network, which is a requirement in this scenario.

Wireless is incorrect because a wireless zone is a broadcasted network connection used within an organization. It would not be appropriate for providing customer access to a product over the internet while also protecting the internal network.

q_attk_deception_dns_sinkhole_secp8

Which of the following is a limitation of using a DNS sinkhole as a cybersecurity measure?

Answers:

- DNS sinkholes can only provide false information for DNS queries.
- ***DNS sinkholes are ineffective if the malware uses a public DNS server or its own DNS server.**
- DNS sinkholes can block all types of malicious traffic.
- DNS sinkholes can prevent malware execution.

Explanation:

If the malware uses a public DNS server or its own DNS server, a DNS sinkhole will not be effective because it won't be able to intercept and redirect the DNS queries.

Providing false information for DNS queries is not a limitation but a function of DNS sinkholes.

While DNS sinkholes can block a lot of malicious traffic by redirecting it, they may not be able to block all types of malicious traffic, especially if the malware uses a different DNS server.

DNS sinkholes do not have the capability to prevent malware execution. They are used to redirect malicious traffic, not to stop malware from executing.

q_attk_deception_email_honeypot_secp8

You are the cybersecurity lead at a large corporation. Recently, your organization has been experiencing an increase in SMTP-based attacks such as open relay, DDoS, and spam attacks. You need to devise a strategy to not only mitigate these attacks but also gather information about the attackers' tactics.

Which of the following would be the BEST solution?

Answers:

- Implement a strong firewall and block all SMTP traffic.
- ***Set up an email honeypot designed to attract and trap these types of attacks.**
- Regularly change the email server's IP address to confuse the attackers.
- Shut down the email server until the attacks cease.

Explanation:

In this scenario, an email honeypot is the best solution. It is designed to attract SMTP-based attacks. While it's distracting the attacker, the cybersecurity team can monitor the malicious activity to learn what the attacker is trying to do, which can be used to strengthen the company's defenses.

While implementing a strong firewall is a good security measure, blocking all SMTP traffic would disrupt legitimate email communication, which is not practical for a large corporation.

Regularly changing the email server's IP address would likely cause significant disruption to legitimate email communication and may not effectively deter determined attackers.

Shutting down the email server would disrupt all email communication, which is not practical for a large corporation. Moreover, it does not provide any opportunity to learn about the attackers' tactics.

q_attk_deception_fake_telemetry_secp8

You are a cybersecurity specialist and you have implemented fake telemetry as part of your organization's defense strategy. An attacker has just probed your network.

Which of the following type of information might your fake telemetry system provide to the attacker?

Answers:

- Login credentials of your organization's employees.
- IP address information of your organization's servers.
- ***False credentials or false IP address information.**
- Credit card information of your organization's customers.

Explanation:

Fake telemetry provides false or spoofed data, such as fake credentials or fake IP address information, in response to network probes.

Providing real login credentials of your organization's employees would compromise the security of your organization, not enhance it.

Providing real IP address information of your organization's servers would expose your organization's network infrastructure to the attacker.

Providing real credit card information of your organization's customers would lead to a serious data breach and is not a function of fake telemetry.

q_attk_deception_honeyfiles_secp8

Which of the following statements about honeyfiles are true? (Select two.)

Answers:

- Honeyfiles are designed to provide real data to the attacker.
- ***Honeyfiles are named in a way that makes them attractive to hackers, enticing them to open or execute them.**
- ***Honeyfiles work with network intrusion detection systems (NIDs) and can help prevent false positives.**
- Honeyfiles are used to block all types of malicious traffic.
- Honeyfiles can only be created by system administrators.

Explanation:

The following are true statements about honeyfiles:

- Honeyfiles are named in a way that makes them attractive to hackers, enticing them to open or execute them.
- Honeyfiles work with network intrusion detection systems (NIDs) and can help prevent false positives by providing a controlled environment for detecting malicious activity.

Honeyfiles do not provide real data to the attacker. They are decoy files set up to attract and trap attackers.

While honeyfiles can help detect and track malicious activity, they do not directly block malicious traffic.

While system administrators often create honeyfiles, end users who are properly trained can also create their own honeyfiles with their own alarms and alerts.

q_attk_deception_honeynet_secp8

You are a cybersecurity specialist at a large corporation. Your company has been experiencing an increase in cyber attacks recently. To better understand the tactics and techniques of the attackers, you have decided to set up a honeynet.

Which of the following is the BEST way to set up and use a honeynet?

Answers:

- Set up the honeynet with real data and systems to make it more attractive to attackers.
- ***Set up the honeynet with decoy systems and monitor it for attacker activity.**
- Set up the honeynet with decoy systems and ignore it until an attack occurs.
- Set up the honeynet with real systems and ignore it until an attack occurs.

Explanation:

A honeynet should be set up with decoy systems that mimic real systems. Monitoring the honeynet allows the cybersecurity team to gather information about the attackers' tactics and tools, which can be used to strengthen the company's defenses.

Using real data and systems in a honeynet is incorrect as it could lead to real data breaches and system compromises. The purpose of a honeynet is to attract attackers with decoy systems, not real ones.

Simply setting up a honeynet and ignoring it until an attack occurs is incorrect because it does not take full advantage of the honeynet's potential. Monitoring the honeynet allows for the collection of valuable information about the attackers.

Using real systems in a honeynet is incorrect because it could lead to real system compromises. The purpose of a honeynet is to attract attackers with decoy systems, not real ones. Additionally, ignoring the honeynet until an attack occurs does not allow for the proactive gathering of information about the attackers.

5.3 Screened Subnets

As you study this section, answer the following questions:

- How is a honeypot used to increase network security?
- What is the typical configuration for a screened subnet?
- What is the function of each firewall in a two-firewall screened subnet?
- What type of computer might exist inside a screened subnet?
- What makes bastion hosts vulnerable to attack? How can you harden bastion hosts?

In this section, you will learn to:

- Configure a screened subnet.

The key terms for this section include:

Term	Definition
Screened subnet	A buffer network (or subnet) that is located between a private network and an untrusted network, such as the internet.
Bastion or sacrificial host	Any host that is exposed to attack and has been hardened or fortified against attack.
Screening router	The router that is most external to the network and closest to the internet.
Duel-homed gateway	A firewall device that typically has three network interfaces. One interface connects to the internet, one interface connects to the public subnet, and one interface connects to the private network.
Screened host gateway	A device residing within the screened subnet that requires users to authenticate in order to access resources within the screened subnet or the intranet.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.5 Given a scenario, modify enterprise capabilities to enhance security. • Firewall ○ Screened subnets
TestOut Security Pro	2.1 Harden Physical Access 2.1.4 Create and configure a screened subnet

5.3.1 Screened Subnets (Lesson Video)

Transcript:

In this lesson, we'll discuss screened subnets. A screened subnet allows you to expose an organization's services to the internet or other untrusted network while protecting the internal network. Implementing a screened subnet is part of a layered security approach.

Let's look at how a screened subnet works. Let's say you manage networks for a small company that provides driver downloads for their products. These drivers are accessed via a customer portal on your website, which is hosted on an in-house web server. The server resides in the company's internal network.
For customers to access this web server and download drivers, this server needs to be exposed to the internet. One way to do this is to place the web server outside the protected network. In other words, on the other side of the perimeter

firewall, right here. This would allow users to access the drivers that reside on this web server. However, this implementation would create a significant security issue.
This is because once you place the server outside the firewall, you have no control over what happens to it. Every hacker on the internet would be able to access this web server very easily and run all types of exploits on it. As you can see, this isn't a very good idea.
However, moving the web server inside the firewall presents another issue. Chances are this firewall allows outbound traffic but blocks all inbound traffic. This makes the web server invisible to the internet, and customers can't access the drivers. To solve this, you could add a rule in the firewall's ACL to allow connections using port 80, HTTP, and port 443, HTTP.
With these ports open, users can initiate connections from the internet to the web server. But, again, this has some risks. You've now opened two significant holes in the firewall here, which has the potential for security breaches. Many exploits take advantage of this scenario.
An attacker gains access to the web server box through the firewall using these two open ports. The attacker then perpetrates a variety of exploits on the web server, and since this web server is on the same network as the rest of the network, the attacker could compromise the entire network. So, what can be done?
One solution is to divide the network into multiple zones of different levels of security. You can create a high-security area and another area within the network that has a lower degree of security. You do this by creating a screened subnet. Let's look at how to do this.
First, you'll need to open ports 80 and 443 on the firewall to let internet users initiate connections to and access the web server. We still have our firewall in place. Opening these ports creates an area of low security on the network. The web server with open 80 and 443 ports is in the low-security area directly behind the firewall. To maintain the security of the rest of your network, you need to install a second firewall between this area of low security and your internal network. By doing this, you create a zone of high security for your internal network.
Because this web server resides in the low-security area with two ports open in the firewall, the web server resides in the screened subnet. The general rule is that you don't put anything in the screened subnet that doesn't absolutely need to be there. For example, the screened subnet isn't the place to put your payroll or your R&D servers. Only the servers and information accessed by users on the internet should be in the screened subnet.
There's still a risk in having this information on the web server. Someone could break in and, for example, replace one of the driver downloads with a modified version with a virus. Be aware of that risk and monitor the screened subnet for it.
The second high-security internal firewall behind the firewall keeps the low-security area of the network to this single segment. The internal network is secure behind the second firewall. When planning your firewall rules, a general rule is that you allow traffic originating in the secured internal network into the screened subnet and through to the internet.
For example, a user on your internal network should be able to open a web browser and access the internet through the various firewalls. However, you must not allow traffic that originates in the low-security area or the no-security area, which is the internet, to establish a connection with a host in the high-security area.
You can also create a screened subnet using a single firewall instead of two. In this scenario, one firewall now has two network connections and two network interfaces installed.
One interface connects to the internal high-security network. The second connects to a separate isolated network segment. You configure this firewall with different sets of rules for the different network segments.
You establish high-security firewall rules for the internal network, creating the area of high security. You implement lower-security rules, such as allowing traffic through ports 80 and 443 for this network segment where the web server resides. This creates a screened subnet security zone.
The lower-security rules for this network segment allow internet connections to the web server on ports 80 and 443. You should harden this server as much as possible because it's sitting in a lower-security area inside the screened subnet.
Each of these options has its strengths and weaknesses. The two-firewall system requires more hardware and, therefore, requires you to administer two separate firewall systems. The benefit is that an attack on the firewall has no effect on the internal firewall. Everything on the high-security network in the high-security zone could be functioning just fine while the firewall is experiencing continual attacks.
The single-firewall solution with multiple interfaces requires less hardware than the first solution. Therefore, you maintain only a single system. The single-firewall solution allows you to create multiple screened subnets. You can create multiple security zones with the one firewall device by adding additional interfaces. You could have a very high-security zone, a medium-security zone, and a low-security zone. The drawback to this approach is that it introduces a single point of failure. If that firewall goes down, everything else goes down with it.
In this lesson, we discussed screened subnets. We first defined what a screened subnet is. We then looked at different ways to create a screened subnet with multiple firewalls or a single firewall with multiple interfaces installed.

5.3.2 Configuring a Screened Subnet (Demo Video)

Transcript:

In this demonstration, we're going to create a screened subnet on our pfSense security appliance. Be aware that screened subnets are also referred to as demilitarized zones, or DMZs. Often, you will see them referred to as DMZs in software settings with various vendors. So, if we encounter something listed as DMZ, we are referring to a screened subnet.

For our current configuration, our device has two connections: LAN and WAN. As mentioned earlier, we are here to configure a screened subnet. I have a few extra optional ports on my pfSense device, and I'm going to configure one of these for the screened subnet. I'll also assign it the IP address 172.16.1.1. Let's get started.
The first thing I'll do from our pfSense Dashboard is confirm our current situation. You can see here that I have WAN and LAN interfaces. To add my screened subnet, I'll go to Interfaces > Assignments. On the Interface Assignments page, I'll click the Add button.
Here, I can see that OPT1 is displayed. I'll go ahead and click on it. OPT1 isn't a very descriptive name, so let's call it 'DMZ.' Before I forget, I'll check the box here to enable the interface. For IPv4 Configuration Type, I'll choose Static IPv4 from the list. We won't configure IPv6 for this demo. Now, let's scroll down to our Static IPv4 Configuration. As per my diagram, I said I was going to assign it the IP address '172.16.1.1,' so that's what I'll type in. I want to give it a /24 subnet mask.
That's all we need to do here. Click Save.
Now we can see the changes in our interface. It also says, "Don't forget to adjust the DHCP Server range if needed after applying." Let's go ahead and click on Apply Changes.
Let's configure our DHCP Server for our screened subnet. I'll go to Services > DHCP Server. We land on the page for the LAN DHCP Server. Let's go to the DHCP Server for the screened subnet.
I'll go ahead and enable DHCP for our screened subnet interface. Now we need to configure the range of IP addresses that our DHCP will hand out. I always leave IPs at the beginning of the range open for static addresses. I can pick something ending in .2 all the way up to .254. I'll set it to '172.16.1.100' and end at '172.16.1.199'. This way, I'll know that anything in the 100s was assigned via DHCP.
I'll scroll down and click Save.
Now let's go back to our dashboard and verify that our screened subnet interface is showing up. I'll scroll down, and right here, you can see that we have our screened subnet. This is where we can add devices to our screened subnet and configure our firewall for handling traffic going to and from the screened subnet.
That's it for this demo. In this demo, we configured a screened subnet.

5.3.3 Configure a Screened Subnet (Simulation)

Scenario

You are the IT administrator for a small corporate network. You want to make a web server that runs services accessible from the internet. To help protect your company, you want to place this server and other devices in a demilitarized zone (DMZ). This DMZ and server need to be protected by the pfSense Security Gateway Appliance (pfSense). Since a few of the other devices in the DMZ require an IP address, you have also decided to enable DHCP on the DMZ network.

In this lab, your task is to perform the following:

- Access the pfSense management console:
  - Username: **admin**
  - Password: **P@ssw0rd** (zero)
- Add a new pfSense interface that can be used for the DMZ.
  - Name the interface **DMZ** .
  - Use a static IPv4 address of **172.16.1.1/16** .
- Add a firewall rule for the DMZ interface that allows all traffic from the DMZ.

    - Use a description of **Allow DMZ to any rule** .
- Configure and enable the DHCP server for the DMZ interface.
    - Use a range of **172.16.1.100** to **172.16.1.200** .

Explanation

Complete this lab as follows:

1. Sign in to the pfSense management console.
    a. In the Username field, enter **admin** .
    b. In the Password field, enter **P@ssw0rd** (zero).
    c. Select **SIGN IN** or press **Enter** .
2. Configure an interface for the DMZ.
    a. From the pfSense menu bar, select **Interfaces** > **Assignments** .
    b. Select **Add** .
    c. Select **OPT1** .
    d. Select **Enable interface** .
    e. Change the Description field to **DMZ** .
    f. Under General Configuration, use the IPv4 Configuration Type drop-down menu to select **Static IPv4** .
    g. Under Static IPv4 Configuration, in the *IPv4 Address* field, enter **172.16.1.1** .
    h. Use the subnet mask drop-down menu to select **16** .
    i. Select **Save** .
    j. Select **Apply Changes** .
    k. (Optional) Verify the change as follows:
        - From the menu bar, select **pfsense COMMUNITY EDITION** .
        - Under Interfaces, verify that the DMZ is shown with the correct IP address.
3. Add a firewall rule to the DMZ interface.
    a. From the pfSense menu bar, select **Firewall** > **Rules** .
    b. Under the Firewall breadcrumb, select **DMZ** . (Notice that no rules have been created.)
    c. Under the Firewall breadcrumb, select **LAN** .
    d. Under the Actions column, select the **copy** icon (two files) for the rule with a source of **LAN net** .
    e. For the Action field, make sure **Pass** is selected.
    f. For the Interface field, use the drop-down menu to select **DMZ** .
    g. For Protocol, make sure it's set to **Any** .
    h. Under Source, use the drop-down menu to select **DMZ net** .
    i. Under Destination, make sure it is configured for **any** .
    j. Under Extra Options, change the description to **Allow DMZ to any rule** . (Is case sensitive.)
    k. Scroll to the bottom and select **Save** .
    l. Select **Apply Changes** .
4. Configure pfSense's DHCP server for the DMZ interface.
    a. From the menu bar, select **Services** > **DHCP Server** .
    b. Under the Services breadcrumb, select **DMZ** .
    c. Select **Enable** .
    d. Configure the Range field as follows:
        - From: **172.16.1.100**
        - To: **172.16.1.200**
    e. Scroll to the bottom and select **Save** .

5.3.4 Screened Subnet Facts

This lesson covers the following topics:

- Screened subnets
- Screened subnet terms

Screened Subnets

A screened subnet, also known as a perimeter network, creates an additional layer of protection between an organization's internal network and the internet. A screened subnet acts as a neutral zone, separating public-facing servers from sensitive internal network resources to reduce the exposure of the internal network resources to external threats. In practical terms, the screened subnet often hosts web, email, DNS, or FTP services. These systems must typically be accessible from the public internet but isolated from sensitive internal systems to limit the impact of a breach of one of these services. By placing these servers in the screened subnet, an organization can limit the damage if these servers are compromised.

Firewalls are typically used to create and control the traffic to and from the screened subnet. The first firewall, between the internet and the screened subnet, is configured to allow traffic to the services hosted in the screened subnet. The second firewall, between the screened subnet and the internal network, is configured to block most (practically all) traffic from the screened subnet to the internal network. A screened subnet is a fundamental part of a network's security architecture and an important example of network segmentation as a type of security control.

Be aware of the following screened subnet facts:

- If the firewall managing traffic into the screened subnet fails, only the servers in the screened subnet are subject to compromise. The LAN is protected by default.
- Packet filters on the firewall allow traffic directed to the public resources inside the screened subnet. Packet filters also prevent unauthorized traffic from reaching the private network.
- When designing the firewall packet filters, a common practice is to close all ports. Open only those ports necessary for accessing the public resources inside the screened subnet.
- To allow access to private resources from the internet, use one of the following approaches:
    - Place a VPN server inside the screened subnet. Require internet users to authenticate to the VPN server and then allow communications from the VPN server to the private network. Only communications coming through the VPN server are allowed through the inner firewall.
    - Copy resources accessible to internet users to servers inside the screened subnet. Even with authentication and authorization configured, this approach exposes those resources in the screened subnet to internet attacks.
- Typically, firewalls allow traffic originating in the secured internal network into the screened subnet and through to the internet. Traffic that originates in the screened subnet (low-security area) or the internet (no-security area) should not be allowed access to the intranet (high-security area).

Only place servers in the screened subnet that need to be there.

Screened Subnet Terms

The following terms are related to screened subnet.

Term	Definition
Bastion or sacrificial host	A bastion host is any host that is exposed to attack and that has been hardened (or fortified) against those attacks. The bastion host is sometimes referred to as a sacrificial host because it is assumed that it will be subject to attack. The term has been applied to the following types of devices:

	• A host that is exposed on the network and is not protected by a firewall device. • The device that provides the firewall service to the screened network behind it. Attacks must pass through the bastion host before they are allowed inside the screened subnet. • A honeypot device that is purposefully exposed to attack in order to distract attackers. The following actions should be taken to harden a bastion host: • Separate roles of bastion hosts by placing a single application on each server. • Fully patch your bastion host on the operating system and on applications. • Run current versions of antivirus and anti-spyware software. • Include a personal firewall. • Uninstall any unnecessary applications or utilities. • Disable and lock down all unnecessary services and ports. • Tighten security on the registry and the user database. • Add IP filters. • Run lockdown facilities, such as IIS lock down and URLScan.
Screening router	A screening router is the router that is most external to your network and closest to the internet. It uses access control lists (ACLs) to filter packets as a form of security. A firewall performing router functions is considered a screening router.
Dual-homed gateway	A dual-homed gateway is a firewall device that typically has three network interfaces: one connected to the internet, one connected to the public subnet, and one connected to the private network. Gateways have to be logged on to, whereas routers pass traffic through without user authentication. IP forwarding is disabled on gateways, effectively blocking through traffic to the network.
Screened-host gateway	A screened-host gateway resides within the screened subnet, requiring users to authenticate in order to access resources within the screened subnet or the intranet.
Two-firewall screened subnet	A screened subnet uses two firewalls. The external firewall is connected to the internet and allows access to public resources. The internal firewall connects the screened subnet to the private network. With a screened subnet, if the outer firewall is compromised, the inner firewall still protects the private network.

5.3.5 Practice Questions (Section Quiz)

q_dmz_bastion_secp8

Which of the following terms describes a network device that is exposed to attacks and has been hardened against those attacks?

Answers:

- ***Bastion or sacrificial host**
- Circuit proxy

- Kernel proxy
- Multi-homed

Explanation:

A bastion or sacrificial host is one that is unprotected by a firewall. The term bastion host is used to describe any device fortified against attack (such as a firewall). A sacrificial host might be a device intentionally exposed to attack, such as a honeypot.

Circuit proxy and kernel proxy are types of firewall devices.

Multi-homed describes a device with multiple network interface cards.

q_dmz_firewall_secp8

You have a company network that is connected to the internet. You want all users to have internet access, but you need to protect your private network and users. You also need to make a web server publicly available to internet users.

Which solution should you use?

Answers:

- ***Use firewalls to create a screened subnet. Place the web server inside the screened subnet and the private network behind the screened subnet.**
- Use firewalls to create a screened subnet. Place the web server and the private network inside the screened subnet.
- Use a single firewall. Put the web server in front of the firewall and the private network behind the firewall.
- Use a single firewall. Put the web server and the private network behind the firewall.

Explanation:

A screened subnet (or DMZ) is a buffer network (or subnet) that sits between the private network and an untrusted network such as the internet. A common configuration uses two firewalls, one connected to the public network and one connected to the private network. Publicly-accessible resources (servers) are placed inside the screened subnet. Examples of publicly-accessible resources include web, FTP, or email servers. Private resources that are not accessible from the internet are placed behind the screened subnet (behind the inner firewall).

Placing the web server inside the private network would mean opening ports in the firewall leading to the private network, which could expose other devices to attack. Placing the web server outside of the firewall would leave it unprotected.

q_dmz_homed_secp8

How many network interfaces does a dual-homed gateway typically have?

Answers:

- 1
- 2
- ***3**
- 4

Explanation:

A dual-homed gateway is a firewall device that typically has three network interfaces: one connected to the internet, one connected to the public subnet, and one connected to the private network.

q_dmz_packets_secp8

What needs to be configured on a firewall to allow traffic directed to the public resource in the screened subnet?

Answers:

- ***Packet filters**
- Subnet
- VPN
- FTP

Explanation:

Packet filters on the firewall allow traffic directed to the public resources inside the screened subnet. Packet filters also prevent unauthorized traffic from reaching the private network.

A subnet is used to segment a network.

A VPN provides a secure outside connection to an internal network's resources. A VPN does not need to be configured on the firewall to allow traffic to the public resource in the screened subnet.

FTP is a protocol used to transfer files. This does not need to be configured on the firewall to allow traffic to the public resource in the screened subnet.

q_dmz_private_secp8

You have used firewalls to create a demilitarized zone. You have a web server that needs to be accessible to internet users. The web server must communicate with a database server for retrieving product, customer, and order information.

How should you place devices on the network to BEST protect the servers? (Select two.)

Answers:

- ***Put the web server inside the screened subnet.**
- ***Put the database server on the private network.**
- Put the database server inside the screened subnet.
- Put the web server on the private network.
- Put the weber server and database server both behind the screened subnet.

Explanation:

Publicly accessible resources (servers) are placed inside the screened subnet. Examples of publicly accessible resources include web, FTP, or email servers. Devices that should not be accessible to public users are placed on the private network.

If you have a public server that communicates with another server, such as a database server, and that server should not have direct contact with public hosts, place the server on the private network and allow only traffic from the public server to cross the inner firewall.

q_dmz_public_secp8

In which of the following situations would you MOST likely implement a screened subnet?

Answers:

- You want internet users to see a single IP address when accessing your company network.
- ***You want to protect a public web server from attack.**
- You want to detect and respond to attacks in real time.
- You want to encrypt data sent between two hosts using the internet.

Explanation:

Use a screened subnet to protect public hosts on the internet, such as a web server, from attack. The screened subnet uses an outer firewall that prevents internet attacks. All publicly-accessible hosts are inside the screened subnet. A second firewall protects the private network from the internet.

Use a Virtual Private Network (VPN) to encrypt data between two hosts on the internet. Use Network Address Translation (NAT) to hide internal IP addresses from the internet. Use an Intrusion Prevention System (IPS) to detect and respond to threats in real time.

q_dmz_screened_subnet_01_secp8

Of the following security zones, which one can serve as a buffer network between a private secured network and the untrusted internet?

Answers:

- Intranet
- Extranet
- Padded cell
- ***Screened subnet**

Explanation:

A screened subnet (or DMZ) is a network placed between a private secured network and the untrusted internet to grant external users access to internally controlled services. The screened subnet serves as a buffer network.

An intranet is a private network that happens to employ internet information services.

An extranet is a division of a private network that is accessible to a limited number of users, such as business partners, suppliers, and certain customers.

A padded cell is an intrusion detection countermeasure used to delay intruders sufficiently to record meaningful information about them for discovery and prosecution.

q_dmz_screened_subnet_02_secp8

Which of the following is the MOST likely to happen if the firewall managing traffic into the screened subnet fails?

Answers:

- All devices in the screened subnet and LAN will be compromised.
- Nothing will happen - all devices will stay protected.

- ***Only the servers in the screened subnet are compromised, but the LAN will stay protected.**
- The LAN is compromised, but the screened subnet stays protected.

Explanation:

If the firewall managing traffic into the screened subnet fails, only the servers in the screened subnet are subject to compromise. The LAN is protected by default.

None of the other options are correct in this scenario.

q_dmz_screen_secp8

Which of the following is another name for a firewall that performs router functions?

Answers:

- Dual-homed gateway
- ***Screening router**
- Screened-host gateway
- Screened subnet

Explanation:

A firewall performing router functions is considered a screening router. A screening router is the router that is most external to your network and closest to the internet. It uses access control lists (ACLs) to filter packets as a form of security.

A dual-homed gateway is a firewall device that typically has three network interfaces: one connected to the internet, one connected to the public subnet, and one connected to the private network.

A screened-host gateway resides within the screened subnet, requiring users to authenticate in order to access resources within the screened subnet or the intranet.

A screened subnet uses two firewalls. The external firewall is connected to the internet and allows access to public resources. The internal firewall connects the screened subnet to the private network.

q_dmz_vpn_secp8

Which of the following is the BEST solution to allow access to private resources from the internet?

Answers:

- Packet filters
- Subnet
- ***VPN**
- FTP

Explanation:

A VPN provides a secure outside connection to an internal network's resources. A VPN server can be placed inside the screened subnet. Internet users can be required to authenticate to the VPN server and then allowed communications from the VPN server to the private network. Only communications coming through the VPN server are allowed through the inner firewall.

Packet filters on the firewall allow traffic directed to a public resource inside the screened subnet. Packet filters also prevent unauthorized traffic from reaching the private network. Packet filters won't allow access to private resources from the internet.

A subnet is used to segment a network.

File Transfer Protocol (FTP) is a protocol used to transfer files. This does not allow access to private resources from the internet.

5.4 Firewalls

As you study this section, answer the following questions:

- What is the difference between a network-based firewall and an application/host-based firewall?
- When would you choose to implement a host-based firewall?
- How are firewall rules used and what are they based on?
- What network security devices can be used for intrusion detection?
- Where should a network-based firewall be placed?

In this section, you will learn to:

- Configure firewall rules.
- Configure firewall schedules.
- Configure a perimeter firewall.

The key terms for this section include:

Term	Definition
Firewall	A device, or software running on a device, that inspects network traffic and allows or blocks traffic based on a set of rules.
Web application firewall (WAF)	A firewall designed specifically to protect software running on web servers and their back-end databases from code injection and DoS attacks.
Network firewall	A firewall that is used to regulate traffic in and out of an entire network.
Stateless firewall	A firewall that allows or denies traffic by examining information in IP packet headers.
Stateful firewall	A firewall that allows or denies traffic based on virtual circuits of sessions. A stateful firewall is also known as a circuit-level proxy or circuit-level gateway.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise
	• Hardening techniques
	  ○ Host-based firewall
	3.2 Given a scenario, apply security principles to secure enterprise infrastructure.
	• Firewall types
	  ○ Web application firewall (WAF)
	  ○ Unified threat management (UTM)
	  ○ Next-generation firewall (NGFW)
	  ○ Layer 4/Layer 7
	4.5 Given a scenario, modify enterprise capabilities to enhance security.
	• Firewall
	  ○ Rules
	  ○ Access lists
	  ○ Ports/protocols

TestOut Security Pro	2.0 Physical and Network Security • Harden physical access ○ Install and configure a firewall

5.4.1 Firewalls (Lesson Video)

Transcript:

In the realm of network security, various tools and technologies stand as sentinels, guarding against the relentless tide of cyber threats. Among these defenders are

Web Application Firewalls, or WAFs, which serve as a defensive wall shielding web servers and their databases from the menacing attempts of code injection and denial-of-service attacks.
Imagine a fortress protecting the heart of a kingdom, the web server, and its valuable database. This fortress, equipped with a WAF, employs a sophisticated set of rules that are application-aware.' This means it not only understands the language of web traffic but also has the ability to detect and respond to threats specific to web applications. It's like having vigilant guards who can identify suspicious behavior in the crowd.
The WAF isn't just a passive guardian; it can be programmed with knowledge of known attack signatures. It acts like an ancient tome, filled with the wisdom of battles past, enabling it to recognize patterns that indicate malicious intent. When it detects these ominous signs, it raises its shield to block the attack. All of this is meticulously documented in logs, a record of the battles fought and threats thwarted, invaluable for understanding and fortifying the defenses.
A WAF can be likened to a powerful talisman, one that can be deployed as an independent appliance, creating a protective zone around the web server.
Alternatively, it can be seamlessly integrated as plug-in software, weaving its protective spells directly into the web server's fabric.
Moving forward, we encounter the Next-Generation Firewall, or NGFW, a transformative figure in the world of network security. The NGFW combines traditional firewall functionalities with advanced capabilities. It performs deep packet inspections, intrusion prevention, and application awareness.

At its core, the NGFW operates with a heightened sensitivity at Layer 7, the application layer. It's like having a security detail that not only checks the ID but also looks into the intentions of each guest. This includes inspecting encrypted traffic.
Moreover, NGFWs extend their protection by seamlessly integrating with network directories, enabling them to assign privileges and restrictions on a per-user or per-role basis. This way, they're equipped to deal with the insider threat, those who bear the disguise of trustworthiness.
Unified Threat Management, or UTM, is another noteworthy character in the realm of network security. It's like a commander who rallies troops from different battalions, bringing them together under one banner.
A UTM centralizes various security controls, like firewalls, anti-malware, and intrusion prevention, into a single, unified appliance. However, this consolidation, while efficient, carries a potential weakness. If the UTM falters, the entire defense line may crumble. It's akin to having all your eggs in one basket.
Additionally, UTM systems may struggle with latency issues bogged down by the weight of too much network activity. They might not perform as well as specialized devices in certain situations.
To some extent, the choice between NGFW and UTM is a matter of scale and need. UTM shines in small and medium-sized businesses where comprehensive security is required, but resources and expertise are limited. NGFW, on the other hand, is the choice for large businesses, offering better performance but with fewer all-encompassing features.
In the world of firewalls, there are two distinguished layers: Layer 4 and Layer 7. Layer 4, akin to a vigilant gatekeeper, inspects the transport layer, the OSI Layer 4. It carefully observes the TCP three-way handshake, ensuring connections are established following the proper protocol. Any deviations, such as suspicious SYN requests or irregular sequence numbers, are swiftly dealt with, preventing malicious attempts to flood or hijack sessions. It can even extend its gaze to UDP traffic, though this is a trickier endeavor due to the inherently connectionless nature of UDP.
Meanwhile, Layer 7, like an astute detective, dives even deeper. It delves into the application layer, the OSI Layer 7, analyzing not only the protocol but also the payload of application-layer packets. This heightened awareness allows it to verify whether the application protocol matches the assigned port.
Imagine a seasoned inspector who checks not only your credentials but also scrutinizes what you're wearing and your behavior.
For example, a Layer 7 firewall can spot a web application trying to send raw TCP data over an HTTP port, revealing hidden threats.
Access control lists, or ACLs, and firewall rules are akin to the rules and regulations governing entry and behavior within a secure facility. Think of ACLs as the bouncers at a club, deciding who gets in and who doesn't. They're like a list of permissions associated with devices like routers and switches.
They work at a network interface level, like a gatekeeper for each entrance.
ACLs use information from incoming packets, like where they're coming from (source IP), where they're going (destination IP), which door they're knocking on (port number), and even what they're saying (protocol).
These lists are used to control traffic across the network.
Now, imagine firewall rules as the instructions given to the bouncers. They tell them how to handle the guests.
Firewalls are like the security system for the whole club (your network). They protect it from unwanted guests and keep the party safe. These rules can be based on various factors, like where the guest is coming from (IP address), which doors they can use (port numbers), what language they're speaking (protocols), or even what kind of dance moves they're showing (application traffic patterns).
Here's the order in which the bouncers (ACLs) check the rules: 1) They start at the top of the list with the most specific rules and work their way down; 2) If a guest matches one of the rules (like having the right VIP pass), they're allowed inside; 3) If none of the rules match, there's usually a final rule that says, "Nobody else is allowed in" (implicit deny); and 4) Some firewalls might not have this final rule by default, so you can add a rule that explicitly denies all access.
It's not just about keeping threats out; firewall rules can also restrict outgoing traffic. Imagine a sieve that filters out undesirable elements from leaving the secure space. For example, a rule can be set to block all outgoing traffic on port 25 (SMTP), preventing a compromised machine from sending out spam emails.
That's it for this lesson. The world of network security is complex and ever-evolving, filled with guardians, gatekeepers, and detectives. Web Application Firewalls, Next-Generation Firewalls, Layer 4 and Layer 7 Firewalls, Access control lists, and firewall rules all play essential roles in fortifying the digital realm. In this digital age, where threats loom at every corner, the knowledge and implementation of proper security measures are paramount to safeguarding data and information.

5.4.2 Firewall Facts

This lesson covers the following topics:

- Firewall types
- ACLs and firewall rules

Firewall Types

A firewall is a network security device or software application that is designed to monitor, filter, and control incoming and outgoing network traffic based on predetermined security rules. The primary purpose of a firewall is to establish a barrier between a trusted internal network (such as a company's local area network or LAN) and untrusted external networks (such as the internet), thereby protecting the internal network from unauthorized access, cyberattacks, and other security threats.

The evolution of network security has given rise to specialized firewalls that offer different levels of protection and control. Depending on the specific security requirements and network architecture, organizations may use a combination of the following firewall types to secure their systems and data.

Type	Description
Host-based firewall	A host-based firewall inspects traffic received by a host. Use a host-based firewall to protect against attacks when there is no network-based firewall, such as when you connect to the internet from a public location. Host-based firewalls are typically software programs. A host-based firewall can be configured to meet the security requirements of the specific host and add an additional layer of security even when a network firewall has been implemented.
Network-based firewall	A network-based firewall inspects traffic as it flows between networks. For example, you can install a network-based firewall on the edge of your private network that connects to the internet to protect against attacks from internet hosts. Network-based firewalls are typically dedicated hardware devices.
Web application firewall	A web application firewall (WAF) is designed to protect software running on web servers and their back-end databases from code injection and denial-of-service attacks. WAFs use application-aware processing rules to filter traffic and perform application-specific intrusion detection. The WAF can be programmed with signatures of known attacks and use pattern matching to block requests containing suspect code. The output from a WAF will be written to a log, which can reveal potential threats to the web application.

Type	Description
	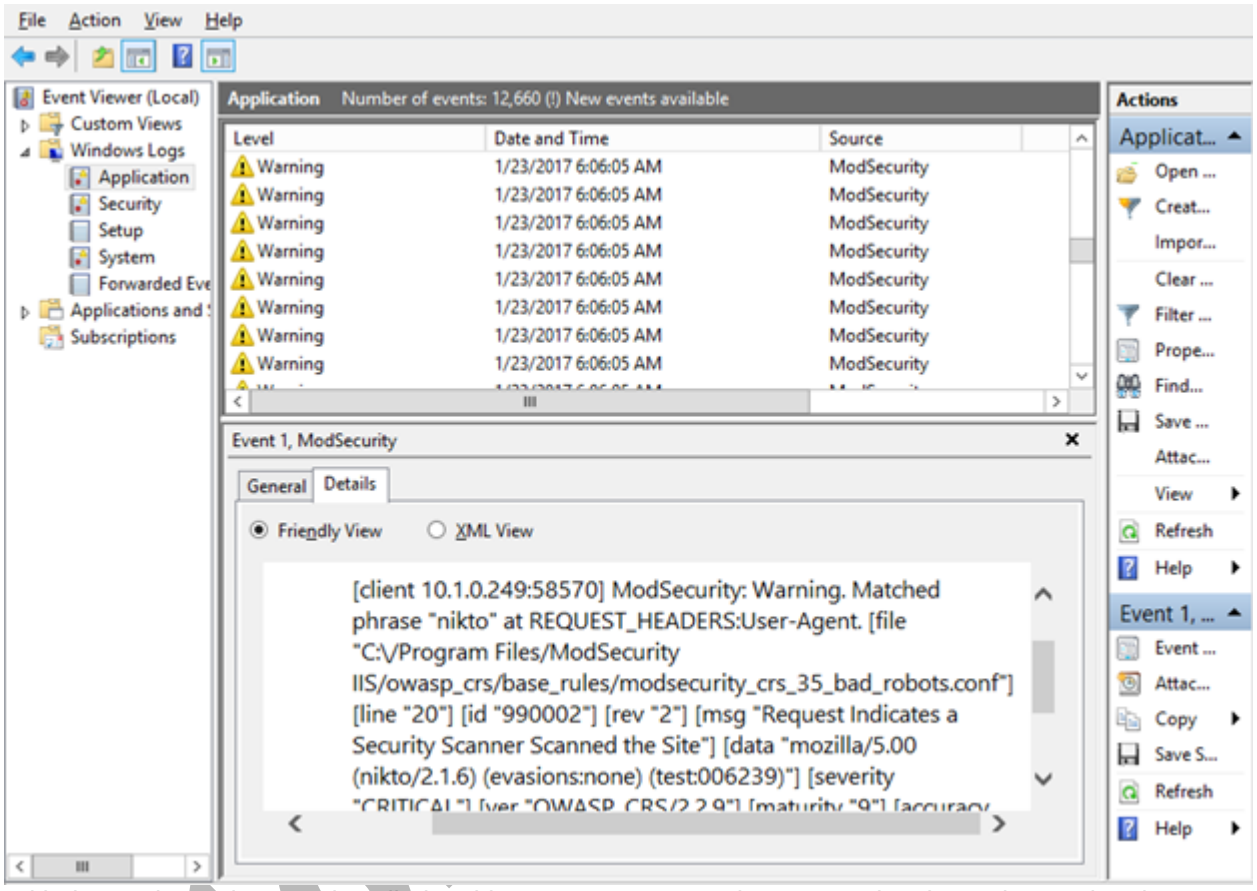 With the ModSecurity WAF installed to this IIS server, a scanning attempt has been detected and logged as an Application event. As you can see, the default ruleset generates a lot of events. (Screenshot used with permission from Microsoft.) A WAF may be deployed as an appliance protecting the zone that the web server is placed in or as plug-in software for a web server platform.
Next-generation firewall	While intrusion detection was originally produced as stand-alone software or appliances, its functionality quickly became incorporated into a new generation of firewalls. The original next-generation firewall (NGFW) was released in 2010 by Palo Alto. There is no official specification for what an NGFW can do, but the following features are typical: • Layer 7 application-aware filtering, including inspection of Transport Layer Security (TLS) encrypted traffic. • Integration with network directories, facilitating per-user or per-role content and time-based filtering policies, providing better protection against an insider threat. • Intrusion prevention system (IPS) functionality. Next-generation firewalls can combine traditional firewall functionalities with advanced capabilities, such as deep packet inspection, intrusion prevention, and application awareness. • Integration with cloud networking.

Type	Description
Unified threat management	Unified threat management (UTM) refers to a security product that centralizes many types of security controls—firewall, antimalware, network intrusion prevention, spam filtering, content filtering, data loss prevention, virtual private networking (VPN), cloud access gateway, and endpoint protection/malware scanning—into a single appliance. This means monitoring and management of diverse controls are consolidated into a single console. UTM has some downsides. • When a defense is unified under a single system, this creates the potential for a single point of failure that could affect an entire network. Distinct security systems, if they fail, might only compromise that particular avenue of attack. • Additionally, UTM systems can struggle with latency issues if subject to too much network activity. Also, a UTM might not perform as well as software or a device with a single dedicated security function. To some extent, NGFW and UTM are just marketing terms. UTM is commonly deployed in small and medium-sized businesses that require a comprehensive security solution but have limited resources and IT expertise. A UTM is seen as a turnkey "do everything" solution, while a NGFW is an enterprise product with fewer features but better performance.
Stateless firewall	A basic packet filtering firewall is stateless. This means that it does not preserve information about network sessions. Each packet is analyzed independently, with no record of previously processed packets. This type of filtering requires the least processing effort, but it can be vulnerable to attacks spread over a sequence of packets. A stateless firewall can also introduce problems in traffic flow, especially when using some sort of load balancing or when clients or servers need to use dynamically assigned ports.
Stateful firewall	A stateful inspection firewall tracks information about the session established between two hosts. All firewalls now incorporate some level of stateful inspection capability. Session data is stored in a state table. When a packet arrives, the firewall checks it to confirm whether it belongs to an existing connection. If it does not, it applies the ordinary packet filtering rules to determine whether to allow it. Once the connection has been allowed, the firewall usually allows traffic to pass unmonitored in order to conserve processing effort.
Layer 4 firewall	Layer 4 is the OSI transport layer. A layer 4 firewall examines the TCP three-way handshake to distinguish new from established connections. A legitimate TCP connection should follow a SYN > SYN/ACK > ACK sequence to establish a session, which is then tracked using sequence numbers. Deviations from this, such as SYN without ACK or sequence number anomalies, can be dropped as malicious flooding or session hijacking attempts. The firewall can be configured to respond to such attacks by blocking source IP addresses and throttling sessions. It can also track UDP traffic, though this is harder as UDP is a protocol without connections. It is also likely to be able to detect IP header and ICMP anomalies.
Layer 7 firewall	A layer 7 firewall can inspect the headers and payload of application-layer packets. One key feature is to verify the application protocol matches the port because malware can try to send raw TCP data over port 80 just because port 80 is open, for instance. As another example, a web application firewall could analyze the HTTP headers and the webpage formatting code present in HTTP packets to identify strings that match a pattern in its threat database. Application-aware firewalls have many different names, including application layer gateway, stateful multilayer inspection, and deep packet inspection. Application-aware devices have to be configured with separate filters for each type of traffic (HTTP and HTTPS, SMTP/POP/IMAP, FTP, and so on).

ACLs and Firewall Rules

An access control list (ACL) is a list of permissions associated with a network object, such as a router or switch, that controls traffic at a network interface level. ACLs typically use packet information like source and destination IP addresses, port numbers, and the protocol to decide whether to permit or deny the traffic. They are usually implemented on network devices to provide traffic control across the network, adding a layer of security and efficiency. In contrast, a firewall rule dictates how firewalls should handle inbound or outbound network traffic for specific IP addresses, IP ranges, or network interfaces. Firewalls typically provide both network and application-level control. They are designed to protect a network perimeter by preventing unauthorized access to or from a network. Firewall rules can be based on various factors, such as IP addresses, port numbers, protocols, or even specific application traffic patterns.

The rules in a firewall's ACL are processed from top to bottom. If traffic matches one of the rules, then it is allowed to pass; consequently, the most specific rules are placed at the top. The final default rule is typically to block any traffic that has not matched a rule (implicit deny). If the firewall does not have a default implicit deny rule, an explicit deny all rule can be added manually to the end of the ACL.

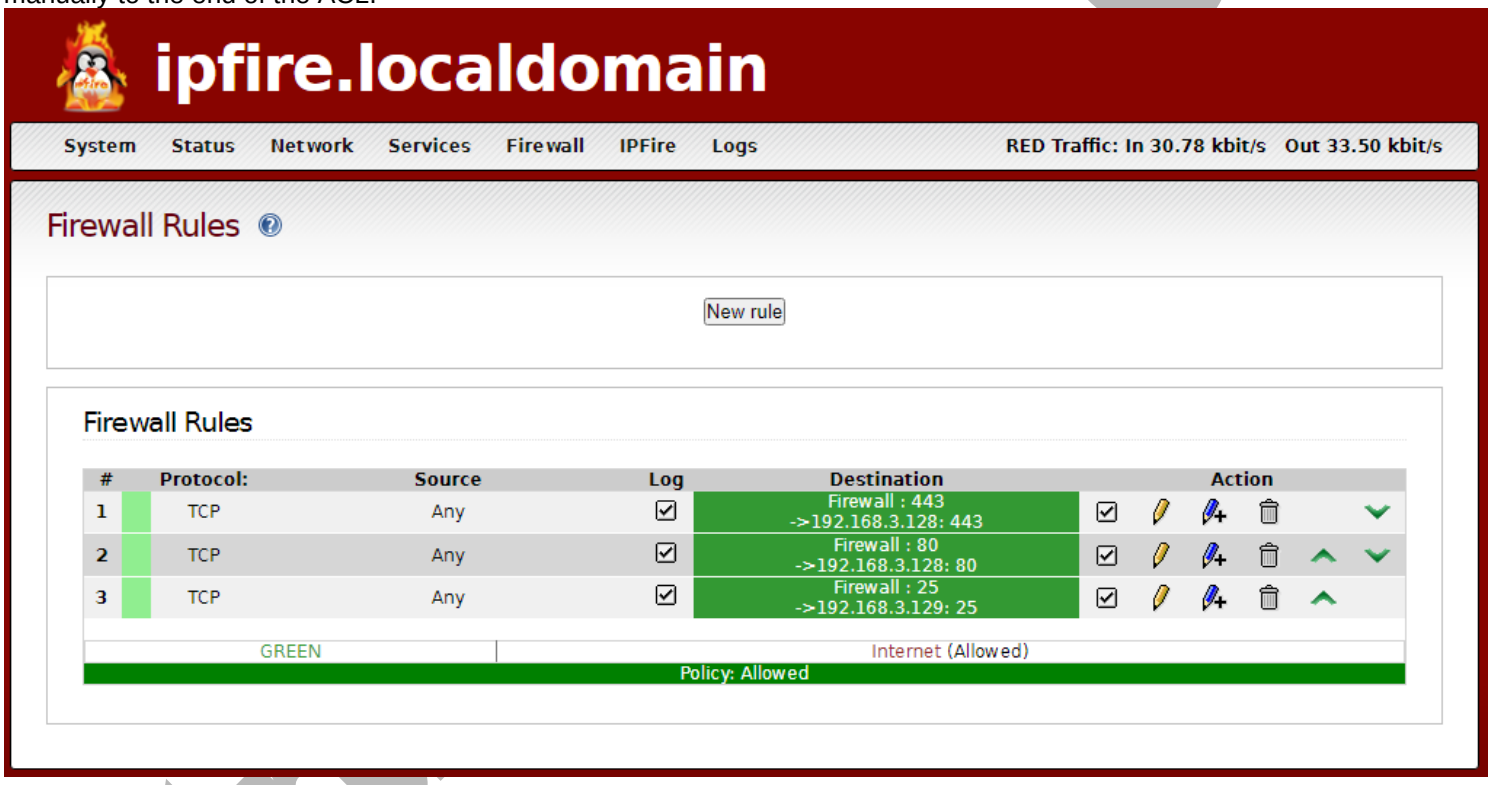

(Screenshot used with permission from IPFire)

Each rule can specify whether to block or allow traffic based on several parameters, often referred to as tuples. If you think of each rule being like a row in a database, the tuples are the columns. For example, in the previous screenshot, the tuples include Protocol, Source (address), (Source) Port, Destination (address), (Destination) Port, and so on.

Even the simplest packet-filtering firewall can be complex to configure securely. It is essential to create a written policy describing what a filter ruleset should do and to test the configuration as far as possible to ensure that the ACLs you have set up work as intended. Also, test and document changes made to ACLs. Some other basic principles include the following:

- Block incoming requests from internal or private IP addresses (that have obviously been spoofed).
- Block incoming requests from protocols that should only function at a local network level, such as ICMP, DHCP, or routing protocol traffic.

- Use penetration testing to confirm the configuration is secure. Log access attempts and monitor the logs for suspicious activity.
- Take the usual steps to secure the hardware on which the firewall is running and use the management interface.

For instance, a firewall rule can be specifically designed to permit or deny traffic based on the TCP or UDP port numbers that a service operates on. If a web server on a network should only allow incoming HTTP and HTTPS traffic, rules could be set up to allow traffic only on ports 80 (HTTP) and 443 (HTTPS), the standard ports for these services. Similarly, rules can be defined to restrict certain protocols such as FTP or SSH from entering the network if they are not needed, thereby reducing the potential attack surface.

Additionally, you can use firewall rules to restrict outgoing traffic to prevent certain types of communication from inside the network. For instance, a rule can block all outgoing traffic to port 25 (SMTP) to prevent a compromised machine within the network from sending out spam emails.

5.4.3 Configuring Firewall Rules (Demo Video)

Transcript:

In this demonstration, we're going to configure firewall rules on a pfSense security appliance. First, we're going to configure a firewall rule to allow both HTTP and HTTPS traffic from the internet, or WAN, through the firewall and to our web server on our screened subnet. The second thing we're going to do is configure a rule to allow all traffic coming from our LAN to get to our screened subnet through the firewall. Be aware that when we are talking about a screened subnet, some vendors may refer to this as a demilitarized zone, or DMZ.

I want to verify that I have my screened subnet ready to go. I'll scroll down, and you can see I have my three interfaces here, and they have assigned IPs.
Let's create our rules. I'll go to Firewall > Rules, and the first rules I'll configure are the ones from our WAN to our screened subnet to allow HTTP and HTTPS to our web server. I want to go to the DMZ tab here, and down here, I have a few buttons. You'll notice there are two buttons that say Add. This one has the arrow pointing up, and if I pick that one, it'll add the rule to the top of my list. If I click on the one with the arrow pointing down, it'll add the rule to the bottom, so I'll click on that one.
For our rule, we want to choose to pass the traffic through the firewall. My other choices are to block it or reject it. If we choose Block, it will just drop the packets as if they never arrived If we choose Reject, the packets are returned to the sender, and the sender can see they were blocked. For security reasons, sometimes you don't want senders to know that your device is even there, and it's better to just block the packets. But if you're troubleshooting issues, reject can be more helpful. Either way, with block or reject, the packets won't reach the destination.
The interface that I want to configure is the DMZ. You can see that I have my LAN and WAN here as well. For Address Family, I'll leave it as IPv4. The other choices are IPv6 or both. For Protocol, we're allowing HTTP, which runs on TCP. If I click the dropdown list, you can see all the other protocols that we could choose if we were going to allow some other type of traffic.
Our source is going to be from our WAN network. I need to see more options, so I'll click on the Display Advanced button. The Source Port Range is going to be HTTP. I could put in a single port or a range, but we just want HTTP port 80. Since this is going to be my web server, I need to put in the IP address here. First, I'll pick Single host or alias, and now let's pop our diagram back up. Our web server's IP is right here, so let's go back and enter that in. Destination Address will be '172.16.1.5'. We're using HTTP, and the rest of this looks good.
Down here, we could log packets handled by this rule. Since this is a web server, we probably don't want to do that since our log would be overwhelmed. If it was FTP, SSH, or other traffic like that, we probably would want to log traffic, but not for normal web requests. We do want to put a description in here because as you create more and more rules, you'll forget what they all do. For this, I'll type in 'HTTP to DMZ from WAN' and click Save. On our next page here, we have to click on Apply Changes for the changes to my firewall to take effect.
So, here are my rules so far. I have one that allows my screened subnet out, and this second one is our HTTP rule from the WAN on port 80 the web server on our screened subnet.
Now, you might be saying to yourself, "Doesn't most web traffic use HTTPS these days?" Well, that answer is certainly yes, so now we need an HTTPS rule.

Here's some good news: to save a bit of time and keep from making any errors, we can just copy this rule and change the HTTP port 80 to HTTPS port 443. I'll click here, which is the Copy icon. What this does creates a new rule with the exact settings. I'll scroll down and change my source from HTTP port 80 to HTTPS port 443. For the destination, I'll do the same. Scroll down change the description by just adding the letter 'S'. Click Save. On this page, Click Apply Changes. Down here, you now see our rules for both HTTP and HTTPS.
Now let's look at our diagram again. This time, I want to create a rule that allows any traffic from my LAN to get to the screened subnet. Let's go back to our firewall and configure that now. These rules are read from the top down. So, for example, if the first rule says block everything, none of the other rules would ever be seen because that's the very first rule. However, if we made a rule right now that said to block everything and put it at the bottom, our first three rules would still be fine, but everything else would be blocked. By default, everything is blocked with pfSense anyway unless you open it, so we wouldn't really need that sort of rule, but it doesn't hurt anything. So far, our rules don't affect one another, so it really doesn't matter where we put them, but keep all of this in mind when creating rules. We'll just click on Add and use the one that puts it at the bottom of the list.
Our Action will be to pass traffic. We'll leave the interface and address family alone. But for protocol, we're going to change this from TCP to Any. This might not be the best practice, but we're allowing any and all traffic from our LAN to reach our screened subnet. Under Source, we'll choose our LAN network from the list. Our destination will be the screened subnet from this list. We don't want to forget a description for this rule. I'll type 'LAN to DMZ Any'. Click Save. As always, Apply Changes.
Here's my latest rule. My source is from the LAN on any port to my DMZ on any port. By the way, the asterisk is wildcard, which means any. Over here, the green check mark means it's enabled.
That's it for this demo. In this demo, we created firewall rules on our firewall.

5.4.4 Configure Firewall Schedules (Demo Video)

Transcript:

There may be times when you want to restrict certain groups of users' access to the internet or other services. For example, you might allow SSH through the firewall to a Linux server, but the person who needs access to it only has it during work hours. Another example is a school that doesn't want normal users accessing the web after hours. There are many cases when you'd want to configure restrictions. For our demo, we're going to use a scenario of blocking web traffic, or HTTP traffic, after business hours.

This is basically a three-part process. First, we'll create a time-restriction schedule. Then we're going to create a rule and apply that schedule to our users. The last thing we'll do is create a rule to override the restriction for members of our management group.
So, under Firewall > Schedules, I go over to Add and click on it. I need to give it a Schedule Name. This can't have any spaces in it, so I'll use underscores and call it Allow_Web_Work _Hours. For our Description we can use anything. I'll call it Allow Web During Work Hours.
For the days, I'll click on the workdays of Monday through Friday to add all of them. For Time, I'll pick 7:00 a.m. from the dropdown list and jump over to Stop Hours. Here, I'll pick 17 and leave the minutes at 59. That makes our stop time 5:59 p.m. I'll click on Add Time, and you can see my Configured Range here, which is now Monday through Friday from 7:00 a.m. to 17:59. That looks good, so let's click Save. Now we're done with our schedule. Next, we need to configure our firewall rules.
We now need to go up to Firewall > Rules and then click on our LAN tab to get to our LAN firewall rules. We have a few rules. One basically allows IPv4 out to anywhere and the same for IPv6.
We want to add our rule to the top of the list. So I'll pick Add with the arrow pointing up.
We want to select Block because we're going to block traffic based on the schedule. I'll scroll down.
We're going to block everything on our LAN network, so I'll pick LAN net from the dropdown list. I need to click on the Display Advanced button here. Now, for our Source Port Range we want to block HTTP port 80.
Let's scroll down a bit. Under Extra Options, click on Display Advance to see additional options. We need to come down and find where it says Schedule. From the list, we have all our schedules. And since we only created one, that's all we have. Let's pick that one. Go down and click on Save.
Now we see our new rule. The source is the LAN network on port 80. Over here, you can see that it refers to the schedule that we created. This all looks good, and so we go up and click Apply Changes.

We have our schedule and we've created a rule that applies that schedule to our users. But what about users that we want to have access to the internet no matter what time of day it is? We need to create another rule for that. But first we need to create an alias. An alias is a group of IPs that'll be grouped together. That alias will be added to a rule that allows internet access.
To do that, we go up to Firewall > Aliases. When the page loads, we go down to the Add button. We need to give it a name that we'll use to identify this group of IP addresses. We'll just say that this group is our management. So, for the Name field that's what I'll put in. I'll do the same for Description.
Now I need to enter in the IP addresses of those devices that'll be part of this alias. I'll enter in a few and start with 10.10.10.1. That IP belongs to Dana. Click Add Host. Now let's enter 10.10.10.2; that'll be Jon's. Click Add Host and we'll add one more, which is 10.10.10.3. That'll be for Mary. Click Add Host one final time and then click Save. Apply the changes and we're ready to create our final rule.
Now that we have our alias, we need to create our rule to override the rule that blocks web access for all users. These rules will be from the top down. If our firewall reads the first rule that says to block all HTTP traffic during non-work hours, that rule is applied. The rule that says to allow management access will never be seen. So, we need to make sure that we tell our firewall that our management can have access before we tell it to block everyone else. We go to Firewall > Rules > LAN and then click the Add button with the up arrow.
All the settings here are fine. But down here, under Source, we need to change this to Single host or alias. Now, watch as I start to type in the address field. As soon as I start to type in Management, it auto-populates with our alias. We'll leave the destination set to Any. For Description, I'll type in Allow Management Web Access Anytime. Click on Save and the rule is created.
As a quick side note, this first rule is created automatically by pfSense. It's there to prevent you from creating any rule that might accidently block you from gaining access to the firewall. That's why it's first in this position. You can actually remove that rule, but that's not part of this lesson. Right below that is our rule. The rule we create for management will be applied before the block web access rule. The one right below it is applied.
That's it for this demo. In this demo, we created a schedule on our firewall to deny web access to our users. Then we created an alias for a group of users who are allowed to bypass that rule with a second rule.

5.4.5 Configure a Perimeter Firewall (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. You recently placed a Web server in the demilitarized zone (DMZ). You need to configure the perimeter firewall on the network security appliance (pfSense) to allow access from the WAN to the Web server in the DMZ using both HTTP and HTTPs. You also want to allow all traffic from the LAN network to the DMZ network.

In this lab, your task is to:

- Access the pfSense management console:
  - Username: **admin**
  - Password: **P@ssw0rd** (zero)
- Create and configure a firewall rule to pass HTTP traffic from the WAN to the Web server in the DMZ.
- Create and configure a firewall rule to pass HTTPS traffic from the WAN to the Web server in the DMZ.
  - Use the following table when creating the HTTP and HTTPS firewall rules:

Parameter	Setting
Source	WAN network
Destination port/service	HTTP (80), HTTPS (443)

Destination	A single host
IP address for host	172.16.1.5
Descriptions	For HTTP: HTTP from WAN to DMZ For HTTPS: HTTPS from WAN to DMZ

- Create and configure a firewall rule to pass all traffic from the LAN network to the DMZ network. Use the description *LAN to DMZ Any* .

Explanation

Complete this lab as follows:

1. Sign in to the pfSense management console.
    a. In the Username field, enter **admin** .
    b. In the Password field, enter **P@ssw0rd** (zero).
    c. Select **SIGN IN** or press **Enter** .
2. Create and configure a firewall rule to pass HTTP traffic from the WAN to the Web server in the DMZ.
    a. From the pfSense menu bar, select **Firewall** > **Rules** .
    b. Under the Firewall breadcrumb, select **DMZ** .
    c. Select **Add** (either one).
    d. Make sure Action is set to **Pass** .
    e. Under Source, use the drop-down to select **WAN net** .
    f. Under Destination, use the Destination drop-down to select **Single host or alias** .
    g. In the Destination Address field, enter **172.16.1.5** .
    h. Using the Destination Port Range drop-down, select **HTTP (80)** .
    i. Under Extra Options, in the Description field, enter **HTTP from WAN to DMZ** .
    j. Select **Save** .
    k. Select **Apply Changes** .
3. Create and configure a firewall rule to pass HTTPS traffic from the WAN to the Web server in the DMZ.
    a. For the rule just created, select the ***Copy*** icon (two files).
    b. Under Destination, change the Destination Port Range to **HTTPS (443)** .
    c. Under Extra Options, change the Description field to **HTTPS from WAN to DMZ** .
    d. Select **Save** .
    e. Select **Apply Changes** .
4. Create and configure a firewall rule to pass all traffic from the LAN network to the DMZ network.
    a. Select **Add** (either one).
    b. Make sure Action is set to **Pass** .
    c. For Protocol, use the drop-down to select **Any** .
    d. Under Source, use the drop-down to select **LAN net** .
    e. Under Destination, use the drop-down to select **DMZ net** .
    f. Under Extra Options, change the Description field to **LAN to DMZ Any** .
    g. Select **Save** .
    h. Select **Apply Changes** .

5.4.6 Practice Questions (Section Quiz)

q_firewalls_acl_01_secp8

Which of the following does a router acting as a firewall use to control which packets are forwarded or dropped?

Answers:

- ***ACL**
- IPsec
- RDP
- VNC
- PPP

Explanation:

When you configure a router as a firewall, you configure the access control list (ACL) with statements that identify traffic characteristics, such as the direction of traffic (inbound or outbound), the source or destination IP address, and the port number. ACL statements include an action to either allow or deny the traffic specified by the ACL statement.

IPsec is a protocol for encrypting packets.

RDP and VNC are remote desktop protocols used for remotely accessing a computer's desktop.

PPP is a protocol for establishing a remote access connection over a dial-up link.

q_firewalls_acl_02_secp8

Which of the following describes how access control lists can be used to improve network security?

Answers:

- ***An access control list filters traffic based on the IP header information, such as source or destination IP address, protocol, or socket number.**
- An access control list looks for patterns of traffic between multiple packets and takes action to stop detected attacks.
- An access control list filters traffic based on the frame header, such as source or destination MAC address.
- An access control list identifies traffic that must use authentication or encryption.

Explanation:

An access control list filters traffic based on the IP header information, such as source or destination IP address, protocol, or socket number. Access control lists are configured on routers, and they operate on Layer 3 information.

Port security is configured on switches, which filter traffic based on the MAC address in the frame.

An intrusion detection system (IDS) or intrusion prevention system (IPS) examines patterns detected across multiple packets.

An IPS can take action when a suspicious pattern of traffic is detected.

q_firewalls_acl_03_secp8

A financial institution has recently undergone a significant network infrastructure upgrade and plans to enhance its network security. It is considering incorporating access control lists (ACLs) into its firewall rules and wants to establish a screened subnet for better security.

Which of the following options would be the MOST beneficial course of action?

Answers:

- ***Incorporate ACLs into its firewall rules and set up a screened subnet**
- Maintain the existing firewall rules and ignore ACLs and screened subnet
- Implement ACLs in its firewall and avoid creating a screened subnet
- Establish a screened subnet but keep the existing firewall rules

Explanation:

Incorporating ACLs into the institution's firewall rules will provide more precise control over traffic based on various parameters. Simultaneously, setting up a screened subnet can isolate public-facing servers from sensitive internal network resources, thus reducing the exposure of the internal network to external threats.

Keeping the current firewall rules without considering ACLs or a screened subnet may leave the institution vulnerable to more advanced threats.

While implementing ACLs in the firewall enhances traffic control and security, not setting up a screened subnet misses an opportunity to reduce the exposure of internal network resources to external threats.

Establishing a screened subnet is a good security practice; however, keeping the existing firewall rules without integrating ACLs misses out on granular traffic control.

q_firewalls_acl_04_secp8

A company's web server is openly accessible to the internet, demanding heightened security measures.

Considering the need for essential protocols and the introduction of a screened subnet, how should the company configure the firewall's access control lists (ACLs)?

Answers:

- ***Permit Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) on firewall access control lists (ACLs); establish a screened subnet for the web server.**
- Block all ports and protocols; enable a screened subnet for the server.
- Allow all ports and protocols; do not create a screened subnet.
- Permit only FTP, SSH, and enable File Integrity Monitoring; disregard screened subnet implementation.

Explanation:

Permitting Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) on the firewall ACLs establishes a screened subnet and adds a layer of security by isolating the server from the internal network, minimizing potential damage from a breach.

Blocking all ports and protocols would hinder the web server's operation. Although a screened subnet would add a layer of security, it doesn't replace the need for well-configured ACLs.

Allowing all ports and protocols introduces unnecessary security risks. Foregoing the creation of a screened subnet misses a crucial opportunity to further isolate and protect the server.

Permitting only FTP and SSH on the firewall ACLs without a screened subnet lacks comprehensive security.

q_firewalls_acl_deny_access_secp8

A security administrator reviews the network configurations of a recently deployed server. The administrator notices that certain unnecessary services have access to the server, potentially creating vulnerabilities. The administrator decides to refine the access control list (ACL) to enhance the server's security.

Which action will the security administrator MOST likely take when refining the ACL to ensure that only necessary services communicate with the server, thereby reducing potential attack vectors?

Answers:

- Permit all incoming traffic to maintain functionality by default
- ***Deny all traffic by default and then allow exceptions based on requirement**
- Permit traffic only from trusted MAC addresses by default
- Implement a stateful firewall for the server

Explanation:

When securing an ACL, best practices dictate denying all traffic by default and permitting only necessary specific traffic. This "deny by default" principle ensures that only approved traffic accesses the resource.

Permitting all incoming traffic creates vulnerabilities since it fails to distinguish between necessary and unnecessary services.

Permitting traffic only from trusted MAC addresses by default does not adhere to the ACL best practice of first denying all traffic by default.

A stateful inspection firewall tracks information about the session established between two hosts, and inspects a packet when it arrives to confirm whether it belongs to an existing connection. If it does not, it applies the ordinary packet filtering rules to determine whether to allow it. This approach does not adhere to the ACL best practice of first denying all traffic by default.

q_firewalls_application_secp8

Which of the following are features of an application-level gateway? (Select two.)

Answers:

- Allows only valid packets within approved sessions
- Uses access control lists
- Verifies that packets are properly sequenced
- ***Reassembles entire messages**
- ***Stops each packet at the firewall for inspection**

Explanation:

Application-level gateways:

- Operate up to OSL Layer 7 (Application layer)
- Stop each packet at the firewall for inspection (no IP forwarding)
- Inspect encrypted packets, such as an SSL inspection
- Examine the entire content that is sent (not just individual packets)
- Understand or interface with the application-layer protocol
- Can filter based on user, group, and data (such as URLs within an HTTP request)

- Is the slowest form of firewall protection because entire messages are reassembled at the Application layer

Allowing only valid packets within approved sessions and verifying that packets are properly sequenced are features of a stateful firewall.

Using access control lists is a feature of a packet-filtering firewall.

q_firewalls_deep_packet_secp8

An IT specialist working for a multinational confectionery company needs to fortify its network security. The firm has been dealing with intrusions where raw User Datagram Protocol (UDP) packets bypass open ports due to a virus.

The specialist will analyze packet data to verify that the application protocol corresponds to the port. The company also wants to track the state of sessions and prevent fraudulent session initiations.

Which of the following tools should the IT specialist prioritize deploying?

Answers:

- ***Deep packet inspection firewall**
- Circuit-level gateway
- Packet filtering firewall
- Transparent firewall

Explanation:

A deep packet inspection firewall acting on layer 7 can scrutinize the content of application packets and confirm protocol-port compatibility. It can also monitor the state of sessions and prevent fraudulent sessions.

A circuit-level gateway, also known as a circuit-level stateful inspection firewall, operates at the session layer and can monitor established connections. However, it falls short of verifying application protocols.

Packet filtering firewalls cannot retain details about network sessions, examining packets independently. They do not maintain session states or validate application protocols.

Although a transparent firewall can review traffic moving between two nodes and operates like a bridge, it cannot scrutinize packet contents at layer 7.

q_firewalls_defense_in_depth_secp8

Which of the following describes the placement and role of a firewall in a network with a defense-in-depth strategy?

Answers:

- A firewall is typically at the network border and serves as a detective control to identify malicious traffic.
- ***A firewall is typically at the network border and serves as a preventive control to enforce access rules for ingress and egress traffic.**
- A firewall is typically inline behind the border firewall and serves as a preventive control to enforce access rules.
- A firewall is typically on internal routers and serves as a corrective control to mitigate denial-of-service (DoS) attacks.

Explanation:

In a network with a defense-in-depth strategy, a firewall is usually at the network border and serves as a preventive control. Its main function is to enforce access rules for traffic entering (ingress) and leaving (egress) the network.

A firewall serves primarily as a preventive control to enforce access rules, not as a detective control to identify malicious traffic.

A firewall is typically at the network border, not inline behind another firewall, to enforce access rules.

While a firewall can mitigate certain types of attacks, it is not typically on internal routers and is not primarily a corrective control for DoS attacks.

q_firewalls_firewall_secp8

Which of the following is the BEST device to deploy to protect your private network from a public untrusted network?

Answers:

- ***Firewall**
- Router
- Hub
- Gateway

Explanation:

A firewall is the best device to deploy to protect your private network from a public untrusted network. Firewalls are used to control traffic entering and leaving your trusted network environment. Firewalls can manage traffic based on source or destination IP address, port number, service protocol, application or service type, user account, and even traffic content.

Routers offer some packet-based access control, but it is not as extensive as that of a full-fledged firewall.

Hubs and gateways are not sufficient for managing the interface between a trusted and an untrusted network.

q_firewalls_hardware_secp8

Jessica needs to set up a firewall to protect her internal network from the internet.

Which of the following would be the BEST type of firewall for her to use?

Answers:

- ***Hardware**
- Tunneling
- Stateful
- Software

Explanation:

Hardware firewalls are physical devices that are usually placed at the junction or gateway between two networks, generally a private network and a public network like the internet. Hardware firewalls can be a standalone product or can also be built into devices like broadband routers.

Software firewalls are generally used to protect individual hosts.

Tunneling is when an attacker wraps a malicious command in an HTTP, ICMP, or ACK tunneling packet that bypasses the firewall and reaches an internal system.

Stateful firewalls, also referred to as stateful multilayer firewalls, determine the legitimacy of traffic based on the state of the connection from which the traffic originated.

q_firewalls_host-based_secp8

A cyber team implements new hardening techniques after a data loss prevention (DLP) audit revealed increased data exfiltration.

What is a tenet of host-based firewalls?

Answers:

- ***It provides controls for incoming and outgoing network traffic.**
- It describes software tools that monitor and protect individual hosts.
- It requires deploying and configuring specialized software agents.
- It uses signature-based detection and anomaly detection.

Explanation:

Host-based firewalls provide controls for incoming and outgoing network traffic and are essential for detecting potential attacks. An important technique for using them when hardening endpoints involves implementing default-deny policies to block all traffic unless explicitly allowed.

Host-based intrusion prevention (HIPS) describes software tools that monitor and protect individual hosts, like computers or servers, from unauthorized access and malicious activities.

HIPS requires deploying and configuring specialized software agents to continuously monitor and analyze endpoints.

HIPS systems use signature-based detection, anomaly detection, and behavior analysis to identify suspicious activities. They also detect and actively respond to threats by automatically blocking or mitigating them.

q_firewalls_host_01_secp8

You have been given a laptop to use for work. You connect the laptop to your company network, use it from home, and use it while traveling.

You want to protect the laptop from internet-based attacks.

Which solution should you use?

Answers:

- ***Host-based firewall**
- Network-based firewall
- VPN concentrator
- Proxy server

Explanation:

A host-based firewall inspects traffic received by a host. Use a host-based firewall to protect against attacks when there is no network-based firewall, such as when you connect to the internet from a public location.

A network-based firewall inspects traffic as it flows between networks. For example, you can install a network-based firewall on the edge of your private network that connects to the internet to protect against attacks from internet hosts.

A VPN concentrator is a device connected to the edge of a private network that is used for remote access VPN connections. Remote clients establish a VPN connection to the VPN concentrator and are granted access to the private network.

A proxy server is an Application-level firewall that acts as an intermediary between a secure private network and the public. Access to the public network from the private network goes through the proxy server.

q_firewalls_host_02_secp8

A new IT tech reviews changes made to a computer, including turning off unneeded logical ports, installing a host-based firewall, and updating the installed antivirus software.

These actions are examples of what technique?

Answers:

- ***Securing the system against a network attack**
- Controlling access to the system
- Implementing host-based intrusion detection/prevention
- Following a baseline configuration and registry settings

Explanation:

Protecting a system from a network-based attack includes disabling unneeded logical ports, installing a host-based firewall, and updating the local antivirus software.

These steps do not assist in controlling access to the system. The steps outlined will help make a system more secure overall, especially when the steps to protect against a network attack pair with steps to harden the system from local attacks.

Host-based intrusion detection and intrusion prevention systems (HIDS/HIPS) require deploying and configuring specialized software agents that continuously monitor and analyze endpoints.

Companies install baseline configuration and registry settings (BCRS) by default and automatically apply the registry settings.

q_firewalls_layer_7_secp8

The network security engineer at a financial corporation is reviewing the current firewall setup. The corporation faces threats from various cyberattacks, some of which leverage application-specific vulnerabilities.

The engineer is considering whether to deploy Layer 4 or Layer 7 firewalls for enhanced security.

If the primary concern is to secure against application-specific attacks, which of the following strategies should the network security engineer consider implementing?

Answers:

- Deploy Layer 4 firewalls on all network edges
- Rely solely on Layer 4 firewalls for internal traffic
- ***Deploy Layer 7 firewalls on all network edges**
- Use Layer 4 firewalls for all internet-facing applications

Explanation:

Layer 7 firewalls, also known as application layer firewalls, have the ability to inspect, control, and often modify application-level data and are effective against application-specific attacks.

Deploying Layer 4 firewalls on all network edges can secure against network-level attacks, but it is not ideal for application-specific attacks as these firewalls mainly operate on the transport layer and do not have deep packet inspection capabilities.

Relying solely on Layer 4 firewalls for internal traffic would miss application-specific vulnerabilities and attacks, as these firewalls do not inspect application-level data.

Layer 4 firewalls are more effective against attacks on the network transport layer, not the application layer.

q_firewalls_ngfw_01_secp8

A network security administrator's responsibilities include enhancing the enterprise's network infrastructure security posture. They deploy a Next Generation Firewall (NGFW) as part of their defense strategy.

The enterprise mixes internal and external services, including a web application and a virtual private network (VPN) for remote access.

Which of the following should the administrator primarily consider when implementing the NGFW to ensure effective security without disrupting normal operations?

Answers:

- ***Deploy the NGFW in inline mode, ensuring it analyzes all traffic while maintaining connectivity.**
- Position the NGFW as a jump server to manage secure access for all network services.
- Set the NGFW to operate in a fail-open mode, ensuring continuous network service even if the firewall fails.
- Use the NGFW as a load balancer, distributing network traffic across multiple servers.

Explanation:

Deploying an NGFW in inline mode enables it to examine all traffic passing through it, identify and mitigate threats, and maintain connectivity without disrupting normal network operations.

An NGFW's primary role is deep packet inspection and threat prevention, not secure access as a jump server.

Although fail-open mode prevents network service interruption if the firewall fails, it may compromise NGFW's primary goal of advanced threat prevention.

A main NGFW function includes advanced threat prevention and deep packet inspection, not load balancing, a technique used to distribute workloads across multiple servers.

q_firewalls_ngfw_02_secp8

A security architect designs a solution to protect the organization's network from advanced threats and provides granular access controls based on user roles.

The organization has a significant volume of TLS-encrypted traffic that needs inspection and wants to integrate the solution with its network directory for role-based content filtering.

Which of the following should the security architect consider the MOST appropriate option?

Answers:

- ***A Next Generation Firewall (NGFW) with Layer 7 application-aware filtering and intrusion prevention system (IPS) functionality**
- A standard stateful firewall with Layer 4 filtering capabilities
- A jump server with enhanced remote access capabilities
- A Web Application Firewall (WAF) designed primarily to protect web applications from targeted attacks

Explanation:

An NGFW with Layer 7 application-aware filtering and IPS functionality can inspect Transport Layer Security (TLS)-encrypted traffic and integrate with network directories for role-based content filtering, meeting the organization's requirements.

A standard stateful firewall with Layer 4 filtering provides basic packet filtering but lacks application-aware filtering at Layer 7 or integration with network directories for role-based content filtering.

A jump server provides secure access to another network segment but lacks advanced threat protection or application-aware filtering required by the organization.

A WAF protects web applications from targeted attacks but does not offer holistic protection or role-based access controls sought by the organization.

q_firewalls_ngfw_utm_secp8

The chief information security officer (CISO) at a medium-sized healthcare company conducts an audit of the company's current security infrastructure.

The company has Next Generation Firewalls (NGFWs) deployed at all external network boundaries, and the CISO is evaluating the possibility of supplementing or replacing the NGFWs with Unified Threat Management (UTM) devices.

If the primary concern is to increase the network's security without introducing significant management complexity, which of the following strategies should the CISO consider implementing?

Answers:

- Replace all NGFWs with UTM devices.
- Deploy UTM devices alongside the NGFWs, with both sets of devices fully active.
- Switch off NGFWs and deploy UTM devices, but only activate UTM when detecting a threat.
- ***Implement UTM devices internally and maintain NGFWs at network boundaries.**

Explanation:

The correct answer is implementing UTM devices internally and maintaining NGFWs at network boundaries. Organizations actively create a balance between enhanced security and manageable complexity when they deploy Unified Threat Management devices internally while also maintaining Next Generation Firewalls at network boundaries.

By replacing NGFWs with UTM devices, organizations actively boost the variety of security features, but this action also introduces significant management complexity due to the increased features that UTM devices have to manage.

Deploying UTM devices alongside the NGFWs, with both sets of devices fully active, does not provide a complete picture of the architecture of implementing UTM devices internally and maintaining NGFWs at network boundaries.

Switching off NGFWs and deploy UTM devices, but only activating UTM devices when detecting a threat does not provide the balance needed between enhanced security and manageable complexity when seeking to increase your network's security.

q_firewalls_packet_01_secp8

Which of the following is a firewall function?

Answers:

- ***Packet filtering**
- FTP hosting
- Encrypting
- Frame filtering

Explanation:

Firewalls often filter packets by checking each packet against a set of administrator-defined criteria. If the packet is not accepted, it is simply dropped.

q_firewalls_packet_02_secp8

You have just installed a packet-filtering firewall on your network.

Which options are you able to set on your firewall? (Select three.)

Answers:

- ***Source address of a packet**
- ***Port number**
- ***Destination address of a packet**
- Sequence number
- Acknowledgement number
- Digital signature
- Checksum

Explanation:

A packet-filtering firewall makes decisions about which network traffic to allow by examining information in the IP packet header, such as source and destination addresses, ports, and service protocols.

Sequence number, acknowledgement number, digital signature, and checksum are information that a packet-filtering firewall does not check by default to make decisions about allowing network traffic through the firewall.

q_firewalls_ports_secp8

When designing a firewall, what is the recommended approach for opening and closing ports?

Answers:

- ***Close all ports; open only ports required by applications inside the network.**
- Open all ports; close ports that expose common network attacks.
- Close all ports; open ports 20, 21, 53, 80, and 443.
- Close all ports.
- Open all ports; close ports that show improper traffic or attacks in progress.

Explanation:

When designing a firewall, the recommended practice is to close all ports and then only open those ports that allow the traffic that you want to allow inside the firewall or the private network.

Ports 20, 21, 53, 80, and 443 are common ports that are opened, but the exact ports you open depends on the services provided inside the firewall.

q_firewalls_stateful_multilayer_secp8

A network administrator at an international baked goods corporation is configuring the company's security infrastructure.

The company has recently had issues with raw Transmission Control Protocol (TCP) packets over open ports by malware, and the administrator needs to be able to inspect packet contents to ensure the application protocol matches the port. The company wants session-state tracking enabled and the ability to block malicious attempts to start bogus sessions.

Which of the following devices should the network administrator focus on implementing?

Answers:

- ***Stateful multilayer inspection firewall**
- Stateful inspection layer 4 firewall
- Stateless packet filtering firewall
- Bridged firewall

Explanation:

A stateful inspection application-aware firewall can inspect the contents of application packets and verify the protocol-port match. It can also track the state of sessions and block bogus sessions.

Layer 4 firewalls operate at the transport layer and can track established connections. However, it cannot verify application protocols like the network administrator requires.

Stateless packet filtering firewalls do not preserve information about network sessions, analyzing packets only in isolation. It does not track session states or verify application protocols.

While a bridged firewall can inspect traffic passing between two nodes and works like a switch, it does not offer the ability to inspect the contents of packets at layer 7.

q_firewalls_stateful_secp8

Which of the following BEST describes a stateful inspection?

Answers:

- Designed to sit between a host and a web server and communicate with the server on behalf of the host.
- ***Determines the legitimacy of traffic based on the state of the connection from which the traffic originated.**
- Offers secure connectivity between many entities and uses encryption to provide an effective defense against sniffing.
- Allows all internal traffic to share a single public IP address when connecting to an outside entity.

Explanation:

Stateful firewalls, also referred to as stateful multilayer firewalls, determine the legitimacy of traffic based on the state of the connection from which the traffic originated. The stateful firewall maintains a state table that tracks the ongoing record of active connections.

A virtual private network (VPN) is a network that provides secure access to a private network through a public network or the internet. Virtual private networks offer secure connectivity between many entities, both internally and remotely. Their use of encryption provides an effective defense against sniffing.

Network Address Translation (NAT) separates IP addresses into two sets. This technology allows all internal traffic to share a single public IP address when connecting to an outside entity.

A firewall can be implemented on circuit-level gateways or Application-level gateways. Both of these firewall designs sit between a host and a web server and communicate with the server on behalf of the host. They can also be used to cache frequently accessed websites for faster web page loading.

q_firewalls_stateless_secp8

Which of the following are characteristics of a basic packet-filtering firewall? (Select two.)

Answers:

- Stateful
- ***Stateless**
- Filters based on URL
- Filters based on sessions
- ***Filters IP address and port**

Explanation:

A packet-filtering firewall makes decisions about which network traffic to allow by examining information in the IP packet header, such as source and destination addresses, ports, and service protocols. A packet-filtering firewall is considered a stateless firewall because it examines each packet and uses rules to accept or reject each packet without considering whether the packet is part of a valid and active session.

A circuit-level proxy or gateway makes decisions about which traffic to allow based on virtual circuits or sessions. A circuit-level proxy is considered a stateful firewall because it keeps track of the state of a session.

Application-level gateways filter on Application layer data, which might include data such as URLs within an HTTP request.

q_firewalls_waf_01_secp8

The security team in a financial organization identified a zero-day vulnerability attack that enables cross-site scripting (XSS) attacks on its internal web portal. The chief information security officer (CISO) instructs the team to take immediate action.

Which action MOST effectively minimizes the threat from the zero-day vulnerability and the potential XSS attacks?

Answers:

- ***Implement a web application firewall (WAF).**
- Upgrade the hardware of the server.
- Encourage staff to change their passwords.
- Restrict the number of login attempts.

Explanation:

Implementing a WAF directly addresses the zero-day vulnerability and XSS attacks by inspecting incoming traffic and blocking suspicious requests.

Upgrading the server hardware might improve overall system performance but does not address software vulnerabilities like XSS.

Encouraging staff to change passwords is generally a good practice in maintaining security, but it does not directly help in preventing XSS attacks. XSS attacks exploit vulnerabilities in web applications to inject malicious scripts, which execute in the context of the victim's session.

Restricting login attempts is a common method used to mitigate brute-force attacks, where an attacker tries multiple combinations to guess a password and does not prevent XSS attacks.

q_firewalls_waf_02_secp8

A newly established e-commerce company experienced increased web-based attacks on its online shopping platform. As a result, the company installed a Web Application Firewall (WAF) to enhance its security infrastructure.

What primary function should the network security manager ensure the WAF is performing to protect the online platform from the MOST common types of web-based threats, such as Cross-site Scripting (XSS), Structured Query Language (SQL) Injection, and Cross-site Request Forgery?

Answers:

- Monitor traffic and block DDoS attacks
- Inspect HTTPS traffic
- ***Validate input and output**
- Encrypt data in transit

Explanation:

A WAF primarily validates input and output. It safeguards against web-based threats by scrutinizing the data sent and received from the web application to ensure compliance with defined security rules.

A WAF assists in mitigating a distributed denial-of-service (DDoS) attack, but its main function does not include handling DDoS attacks. A more comprehensive network security infrastructure manages DDoS attacks.

Though inspecting HTTPS traffic contributes to threat identification, it aligns more with a network-based firewall or an intrusion prevention system function than a WAF.

While encrypting data in transit is crucial to protect sensitive data, it relates more to using protocols such as transport layer security or Secure Sockets Layer than being a primary function of a WAF.

5.5 Virtual Private Networks

As you study this section, answer the following questions:

- What are three ways a Virtual Private Network (VPN) can be implemented?
- What is a VPN concentrator?
- What function do VPN endpoints provide?
- What is the difference between full tunnel and split tunnel?
- What are three types of protocols used by a VPN?
- How does a transport layer security (TLS) VPN work?

In this section, you will learn to:

- Configure a VPN.
- Configure a VPN client.
- Configure a remote access VPN.
- Configure a VPN connection iPad.

The key terms for this section include:

Term	Definition
Virtual Private Network	A remote access connection that uses encryption to securely send data over an untrusted network.
Tunneling	The practice of encapsulating data from one protocol for safe transfer over another network such as the Internet.
Point-to-Point Tunneling Protocol (PPTP)	A early tunneling protocol developed by Cisco and Microsoft to support VPNs over PPP and TCP/IP. PPTP is highly vulnerable to password cracking attacks and considered obsolete.
Layer 2 Forwarding (L2F)	A tunneling protocol developed by Cisco to establish virtual private network connections over the internet.
Internet Protocol Security (IPsec)	Network protocol suite used to secure data through authentication and encryption as the data travels across the network or the Internet.

Secure Sockets Layer (SSL)	A well-established protocol to secure IP protocols, such as HTTP and FTP. And can also be used to secure other application protocols and as a virtual private networking (VPN) solution.
Transport Layer Security (TLS)	Security protocol that uses certificates for authentication and encryption to protect web communications and other application protocols.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
TestOut Security Pro	2.2 Harden Network Devices 2.2.3 Configure and Access a Virtual Private Network (VPN)
CompTIA Security+ SY0-701	3.2 Given a scenario, apply security principles to secure enterprise infrastructure. • Secure communication/access ○ Virtual private network (VPN)

5.5.1 Virtual Private Networks (Lesson Video)

Transcript:

A virtual private network (VPN) uses encryption to send data securely over an untrusted network. The high cost of dedicated WAN connections was one of the mains reasons for VPNs to be used by organizations.

A VPN takes advantage of an existing internet connection to communicate securely between devices. Let's see how this is done.
VPNs provide a secure internet connection between locations by encrypting packets in transit. A VPN uses a protocol that tunnels, or encapsulates, each of those packets into a new packet.
Information in the packet header of these encrypted packets routes the information through the internet. On the destination device, the outer wrapping of the packets is removed, and the packet is decryptedâ€"the data is back in its original format.
If an attacker were to intercept packets sent via a VPN, they wouldn't be able to read any of the contents. Only the device on the other end has the appropriate decryption key that allows them to view the packet contents.
Now, there are multiple ways that you can configure a VPN through the internet.
The first way is using a host-to-host VPN. This type of VPN allows an individual host connected to the internet to establish a VPN connection to another host. With a host-to-host connection, both devices need the ability to establish and understand the VPN protocol that's used. Both devices must have the software for encrypting the packets and encapsulating the packets.
Another way is using a site-to-site VPN. With a site-to-site VPN, you have a collection of computers at each location. Each computer in any location is able to communicate securely with any other computer at another location.
Rather than requiring VPN configuration on every single computer, you install a single device, which acts as a gateway server. A VPN gateway is a dedicated VPN device that handles VPN connections, as well as the encryption and decryption of packets sent over the internet.

The nice thing about using a VPN gateway is the protocols and encapsulation method only need to be configured and maintained on a single device for the entire network, instead of on each individual device. It also allows for all traffic between the two sites to be encrypted 100 percent of the time.
The final method is a remote access VPN. In this case, individual hosts on the network can establish a VPN connection to the remote site. In this configuration, the client computer must be able to establish the VPN connection with a special device called a VPN concentrator that sits on the edge of the private network.
VPN concentrators are hardware devices that are dedicated to establishing client connections, as well as encrypting and decrypting VPN packets. Each client is configured with software that allows it to encrypt packets. The VPN concentrator is configured to allow or reject connections from users. It also removes the encryption before forwarding the information to the private network.
Now, VPNs can use a few different methods to establish connections and encrypt traffic. These methods are known as VPN tunneling protocols. The two most commonly used VPN tunneling protocols are IPsec and SSL.
IPsec encrypts contents sent through a tunnel created by another protocol. IPsec is probably one of the most common tunneling encryption mechanisms currently used, and IPsec is actually made up of two different protocols.
The first one is the Authentication Header, or AH. The second one is the Encapsulating Security Payload (ESP). AH is used to authentic the connection, while ESP is used to encrypt the data that's being sent through the connection.
When you implement IPsec, you can use either of these protocols by themselves. For example, if you chose to implement AH only, then your VPN will have authentication, but no encryption. In most cases, you'll want to make sure that your VPN solution using IPsec includes ESP.
IPsec also has two different modes for sending packets through the tunnel. The first way is called transport mode.
With transport mode, IPsec only encrypts each packet's internal data. Everything elseâ€”the destination IP address, the origination IP addressâ€”is all in cleartext.
The second mode is called tunneling mode. In tunneling mode, the entire packet is encrypted. It is then encapsulated in another non-encrypted packetâ€”complete with a new IP headerâ€”and sent over the internet. Tunnel mode is usually the default sending mode for IPsec.
Another protocol that can be used for a VPN connection is SSL. SSL's been around for a long time. It's been used in combination with other protocols, mostly HTTP, in order to secure traffic between two devices.
For example, if you were to go to an online store and purchase a product, more than likely, you're going to use the secure version of HTTPâ€”HTTPSâ€”to protect the credit card information we're sending to the web server.
Well, HTTPS leverages SSL to encrypt the traffic between those two devices. This makes SSL a great option for encrypting other types of connections between devices, such as a VPN connection. SSL requires certificates for identity proof, as well as for encryption.
One of the benefits of using SSL is the fact that it uses port 443. This is really important because most network firewalls in most organizations are already configured to allow HTTPS traffic on this port, so we don't have to make any major firewall changes if we want to deploy an SSL VPN.
There's one more VPN configuration option we need to discuss, and that is the different between split tunnel and full tunnel VPN configurations.
With a split tunnel VPN, only certain types of trafficâ€”for example, traffic destined for a specific IP address rangeâ€”are sent through the VPN connection. All other traffic goes through the internet as normal. This configuration might be good for people who need to access private network resources but still want to access the internet through their own internet, and not through the VPN. It also helps reduce the amount of traffic sent through the VPNâ€”instead of sending all traffic through the VPN, only necessary traffic is sent.
Split tunneling also has something called inverse split tunneling. This is where all traffic is sent through the VPN except for a specific type of traffic, which is routed through the regular internet, unencrypted. The split is inverted, as its name suggests.
The other way you can configure the VPN is to route all traffic through the VPN, regardless of the type of traffic. This is usually the default VPN configuration method.
When you implement a VPN, be sure to select a protocol supported by all of the devices that need to encrypt or encapsulate packets. When you use a VPN through a firewall, make sure you open the necessary ports to allow the VPN traffic through the firewall. In addition, make sure you know which type of VPN connection and configuration is best for your organization.

5.5.2 Configuring a VPN (Demo Video)

Transcript:

In this demonstration,Â we're going to cover how to configure OpenVPN on a pfSense security appliance.Â We're already logged into pfSense on the Dashboard page.

The first thing I'm going to do is to go to System > Package Manager and select Available Packages. I'm going to type in â€˜openvpn' to search for the OpenVPN Client Export package. This will allow us to export our configuration settings to make setting up our clients a lot easier. Click Install and Confirm. This will take about a minute or so to install. Now let's confirm the installation. Success.
Now let's go over to VPN > OpenVPN. We'll select Wizard to get started. In the past, setting up a VPN was a bit complicated, but the wizard makes setting up OpenVPN really quick and easy.
For the type of server, we're going to select Local User Access. We could also pick LDAP or RADIUS if you have those set up, but we don't. Click Next to continue.
Now we need to create a certificate authority. pfSense adds this right in the wizard so you don't have to do it manually somewhere else, which is nice.
For the Descriptive name, I'll just enter in "TestoutCA". I'll leave these next few fields at the default values and go down to Country Code. Here, I'll enter â€˜US'. For State or Province, I'll enter â€˜UT' for Utah. For City, I'm in the wonderful town of â€˜Pleasant Grove'. For Organization, I'll put â€˜TestOut'. Now click on Add new CA to add the Certificate Authority to the server.
Now we must create a new server certificate. We just added the Certificate Authority. Next, we need an actual certificate. For Descriptive name, I'll just put â€˜Testout'. I'll leave everything with the default values. Country Code, State, City, and Organization are all auto-populated with the correct answers, so I'll just click on Create new certificate.
Now let's configure the server information. For our interface, we want to select our WAN. We also have LAN and DMZ. Our WAN is where traffic will typically come from. Our Protocol will be UDP on IPv4 only. You can see we have some other choices in case your setup is different from mine.
Our VPN uses the default port of 1194. We'll enter a description for our VPN and just call it â€˜TestoutVPN'.
We'll leave all the cryptographic settings set to their defaults and skip down to Tunnel Settings.
The first thing we must do is enter some virtual network settings for our tunnel network. This will be a virtual network that our clients and server will use to communicate on. I'm going to enter '10.10.20.0/24'. All my VPN clients will get an IP address of 10.10.20.something with a subnet mask of 255.255.255.0.
Next is our local network. This is the LAN network that your VPN clients will access. My local LAN network is '10.10.10.0/24'.
For concurrent connections, I'll just put in 5. You can enter more or less than that. It just depends on how many VPN users you'll have connecting at one time.
Now, once your users connect, they'll need a DNS server. You could put in a public DNS or local DNS. I'll put in the DNS that my LAN uses, which is '10.10.10.1'. This is also the IP address of my pfSense appliance.
Scroll down and click Next.
Now we need to set up two firewall rules. We need to allow traffic from our VPN clients to our VPN server. We also need to allow traffic from our clients through the VPN to our network. Lucky for us, we don't have to do this manually. OpenVPN understands what we need, and by checking these boxes, OpenVPN will configure those rules for us. Configuring rules is not part of this lesson, but they're certainly a skill you need to know. Click Next to continue. And on this page, click Finish.
Let's review what we've done. We have our interface set to WAN. Our protocol and port are set to UDP on port 1194. Our tunnel network looks good. Our encryption is here. And our description is TestoutVPN.
Now let's head over to System > User Manager and click on Add. I really don't want my admin user to be the same as one of my VPN users, so I'm going to add a different user here to keep them separate. For the username, I'll enter â€˜danafellows'. I'll enter my password and confirm it. I'll just put Dana for the full name. Click Save.
Everything looks good, and that's all I need to do here.
I'll go back up to VPN > OpenVPN, come down here, and click on the edit icon. For some reason, our server mode changes to Remote Access SSL plus TLS, but we're using Remote Access (User Auth), so we need to confirm that and change it back if needed. It looks like we do, in fact, need to change that. I'll scroll down to the bottom and click on Save.
Now it's time to download the client configuration. We'll go to Client Export. On the next page, scroll down to our list of OpenVPN client configuration options. I'll find the one for the Windows Installer and click on it to download the configuration file.
Now I'll copy this over to a machine I have on a different network and install the OpenVPN client.
I'm now on a machine on my DMZ network. I copied the installation file over, and I'll install it now. I'll click Yes to the User Access Control, click Install to continue, and click through all the rest of these to complete the installation process. Click Finish, and finally, click Close.

Now let's double-click to start the OpenVPN client. I'll drag the installation files into the Recycle Bin.
I'll move the OpenVPN client, here on the taskbar, to a spot where we can see it better. Right-click on it and choose Connect. My username is automatically populated, so I'll just enter in my password and check the Save password box. Click OK to connect. I get a message that I'm connected, and my Assigned IP is listed here.
Now, if you're like me, you might ask yourself, "How did the client know how to connect to my OpenVPN server out there on the internet?" Let's take a second to look at the config file. I'll right-click and select Edit Config to open it. It opens in a text file, and here are all the settings that were configured for us. Down here is the IP address of the WAN interface.
Now, you might say, "Hey, that's a private IP address." Yes it is. This is a test network, so I'm actually using a private IP range for my WAN. Let's take a look at that for a minute.
Here's a diagram of my test network. I accessed my pfSense and OpenVPN settings from this client, over here. The configuration is actually here, on the pfSense. The client over on my DMZ network is the one that initiated the connection to the VPN just a second ago. This client can be anywhere on the internet, but for demo purposes, I placed it on my DMZ. Now let's just confirm a few final things before we wrap up.
I'm still on my client on my DMZ. I'm going to open a command prompt to check my IP settings. I'll do an â€˜ipconfig /all' and press Enter. I'll scroll up a bit. Here's the IP address of this machine on the DMZ, 172.16.1.101. That is my IP scheme for the DMZ. Now let's scroll down a bit. Remember, we made our tunnel network 10.10.20.0, and right here is an address from that range that's assigned to us. We saw that a few minutes ago, also right after we connected.
One final thing. I access the pfSense configuration through a web browser. I have it blocked from being accessed externally. But since we're connected to the LAN through this VPN, let's see if we can access it. I'll open my browser, and in the Address field, I'll enter the IP of the pfSense with the LAN address of 10.10.10.1. When I do that, the Sign In page is displayed, confirming that I have access to LAN resources.
That's it for this demo. We configured OpenVPN on pfSense. We installed the client configuration package, set up our certificate, configured OpenVPN server settings, created a user, downloaded and installed the OpenVPN client on a remote host, and confirmed our connection.

5.5.3 Configuring a VPN Client (Demo Video)

Transcript:

In this demonstration, we're going to create a client VPN connection on a Windows 10 system.

There's VPN software built into Windows 10 that works really well. We're going to go down to the Start button and choose the Settings option. Inside the Settings option, notice that there's a category for Network and Internet, and inside of this category, there's a node for VPN. We can create a VPN connection, and we can choose to allow this VPN connection over metered networks.
This means if we're connecting, maybe via cellular, and this is going to be a charge per gigabyte of data that we transfer, we're going to permit the VPN over this type of connection. We'll also permit the VPN while we're roaming with the system by default, so we'll choose the add a VPN connection button, and then we can choose the VPN provider.
Let's say that we're using the built-in functionality of VPNs with Windows as our VPN provider, and we're going to say this is 'CorpNetVPN'. Then we'll provide the IP address or the fully qualified domain name as long as it's registered in your DNS system for the VPN server. This is typically information that you're going to provide your end users so that they can successfully make the VPN connection.
For the VPN type, often, we'll configure it to be automatic, and we'll have given the instructions to this client for how to connect. But notice you can also configure it so that it's preset for the VPN type. A popular option here is IKE version 2. This would be used with IPSec VPN connections, for example.
You can easily use a smart card, a one-time password, a certificate, or a username and password for verification. We'll do a dfellows username and a password that's extra strong. Once the correct information is entered, we'll save it, and now have a CorpNet VPN connection. We can connect at any time we'd like, or we could go in the advanced options and edit the properties of this connection, clear the sign-in info, or specify that we need to use a VPN proxy. This is a device we'd go to for establishing the VPN on our behalf.
We can easily remove any VPN connection that we've created inside of this settings area, provided that we have the appropriate permissions inside of Windows 10.
That's it for this demo. We discussed how easy it is to establish a VPN connection inside of Windows 10.

5.5.4 Configure a Remote Access VPN (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. Occasionally, you and your co-administrators need to access internal resources when you are away from the office. You would like to set up a Remote Access VPN using pfSense to allow secure access.

In this lab, your task is to use the pfSense wizard to create and configure an OpenVPN Remote Access server using the following guidelines:

- Sign in to pfSense using:
  - Username: admin
  - Password: P@ssw0rd (zero)
- Create a new certificate authority certificate using the following settings:
  - Name: **CorpNet-CA**
  - Country Code: **GB**
  - State: **Cambridgeshire**
  - City: **Woodwalton**
  - Organization: **CorpNet**
- Create a new server certificate using the following settings:
  - Name: **CorpNet**
  - Country Code: **GB**
  - State: **Cambridgeshire**
  - City: **Woodwalton**
- Configure the VPN server using the following settings:
  - Interface: **WAN**
  - Protocol: **UDP on IPv4 only**
  - Description: **CorpNet-VPN**
  - Tunnel network IP: **198.28.20.0/24**
  - Local network IP: **198.28.56.18/24**
  - Concurrent Connections: **4**
  - DNS Server 1: **198.28.56.1**
- Configure the following:
  - A firewall rule
  - An OpenVPN rule
- Set the OpenVPN server just created to **Remote Access (User Auth)** .
- Create and configure the following standard remote VPN users:

Username	Password	Full Name
blindley	L3tM31nNow	Brian Lindley
jphillips	L3tM31nToo	Jacob Phillips

Explanation

While completing this lab, use the following information:

- Create and configure the following standard remote VPN users:

Username	Password	Full Name
blindley	L3tM31nNow	Brian Lindley
jphillips	L3tM31nToo	Jacob Phillips

Complete this lab as follows:

1. Sign in to the pfSense management console.
    a. In the Username field, enter **admin** .
    b. In the Password field, enter **P@ssw0rd** (zero).
    c. Select **SIGN IN** or press **Enter** .
2. Start the VPN wizard and select the authentication backend type.
    a. From the pfSense menu bar, select **VPN > OpenVPN** .
    b. From the breadcrumb, select **Wizards** .
    c. Under *Select an Authentication Backend Type* , make sure **Local User Access** is selected.
    d. Select **Next** .
3. Create a new certificate authority certificate.
    a. For Descriptive Name, enter **CorpNet-CA** .
    b. For Country Code, enter **GB** .
    c. For State, enter **Cambridgeshire** .
    d. For City, enter **Woodwalton** .
    e. For Organization, enter **CorpNet** .
    f. Select **Add new CA** .
4. Create a new server certificate.
    a. For Descriptive Name, enter **CorpNet** .
    b. Verify that all of the previous changes (Country Code, State/Providence, and City) are the same.
    c. Use all other default settings.
    d. Select **Create new Certificate** .
5. Configure the VPN server.
    a. Under General OpenVPN Server Information:
        - Use the *Interface* drop-down menu to select **WAN** .
        - Verify that the Protocol is set to **UDP on IPv4 only** .
        - For Description, enter **CorpNet-VPN** .
    b. Under Tunnel Settings:
        - For Tunnel Network, enter **198.28.20.0/24** .
        - For Local Network, enter **198.28.56.18/24** .
        - For Concurrent Connections, enter **4** .
    c. Under Client Settings, in DNS Server1, enter **198.28.56.1** .
    d. Select **Next** .
6. Configure the firewall rules.
    a. Under *Traffic from clients to server,* select **Firewall Rule** .
    b. Under *Traffic from clients through VPN* , select **OpenVPN rule** .
    c. Select **Next** .
    d. Select **Finish** .
7. Set the OpenVPN server just created to **Remote Access (User Auth)** .
    a. For the WAN interface, select the **Edit Server** icon (pencil).
    b. For Server mode, use the drop-down and select **Remote Access (User Auth)** .
    c. Scroll to the bottom and select **Save** .
8. Configure the following Standard VPN users.
    a. From the pfSense menu bar, select **System > User Manager** .

b. Select **Add** .
c. Configure the User Properties as follows:
    - Username: ***Username***
    - Password: ***Password***
    - Full name: ***Fullname***
d. Scroll to the bottom and select **Save** .
e. Repeat steps 8b-8d to create the remaining VPN users.

5.5.5 Configure a VPN Connection iPad (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. You recently set up the Remote Access VPN feature on your network security appliance to provide you and your fellow administrators with secure access to your network. You are currently at home and would like to connect your iPad to the VPN. Your iPad is connected to your home wireless network.

In this lab, your task is to:

- Add an IPSec VPN connection using the following values:

  This can be added by selecting **Settings** > **General** > **VPN** .

Parameter	Value
Description	**CorpNetVPN**
Server	**198.28.56.22**
Account	**mbrown**
Secret	**asdf1234$**

- Turn on the VPN.
- Verify that a connection is established. The password for mbrown is **L3tM31nN0w** (0 = zero).

Explanation

Complete this lab as follows:

1. Verify your connection to the Home-Wireless network.
    a. Select **Settings** .
    b. Select **Wi-Fi** .
    From the right, notice that you are connected to the Home-Wireless network.
2. Add and configure a VPN.
    a. From the left menu, select **General** .
    b. From the right menu, select **VPN** .
    c. Select **Add VPN Configuration** .

d. Select **IPSec** .
e. Configure the IPSec options as follows:
- Description: **CorpNetVPN** .
- Server: **198.28.56.22**
- Account **: mbrown**
- Secret: **asdf1234$**

f. In the upper right, select **Save** .

3. Connect to the VPN just created.
a. Under VPN Configuration, slide Not Connected to **ON** .
b. When prompted, enter **L3tM31nN0w** (0 = zero) as the password.
c. Select **OK** .

5.5.6 Configure Remote Access, non VPN (Demo Video)

Transcript:

Remote access is a method that allows you to connect to a computer remotely from another computer. This could be useful if you're logging into your computer at home while you're away, helping a family member, or just requiring access to a server. Some of you may be familiar with a VPN, which stands for Virtual Private Network. This is another means of connecting to a remote network when you're physically not there. These connections require more knowledge to set up, but the non-VPN method we're showing you today doesn't.

To get started, we'll open our web browser and type in remotedesktop.google.com. We'd like to access my computer, so let's click on this option here. You'll be required to have a Google account. We already have one, so we're just going to log in. First, we need to provide our email, and after that, input our password. We're not worried about saving the password or this extra pop-up down here, so we'll make these go away. On the left-hand side, we have three options: Remote Access, Remote Support, and Set up via SSH. Remote Access is for accessing another PC, something you'd set up. Remote Support would be for helping someone else, where they'd need to download a link and provide you with a code. The third option is to set up an SSH connection to a Windows or Linux Computer.
We're going to walk through the setup for remote access. First, we must ensure we're on the computer requiring remote access and then download this link here. It will download an extension that can be used by both Chrome and Edge, as they're built on the same platform. It may take a bit for this to download, so we're going to fast-forward. Now, you can see the MSI file is downloaded. All we have to do is click Accept & Install. This will start the installation process. The window should pop up to show us where the MSI file was downloaded. Once executed, it will prompt to close some applications before it can be installed. Our install is done, so we can close this window.
Now that we have Chrome Remote Desktop installed, we can name the PC we are setting up. By default, it grabs the hostname of the computer. If your hostname wasn't set up previously, you may see a mix of letters and numbers instead of Office-PC like this one. We're going to click Next. For security purposes, you'll be required to set up a PIN. I'm going to set this and then confirm it for a second time. After clicking Start, we should be able to accept one final UAC prompt, and then our Office-PC should be ready to go. You'll notice it says Offline; however, when you refresh the web browser, it's now Online. Now, let's try connecting to this PC from another PC.
Okay, we're on our other PC where Chrome Remote Desktop isn't installed. Let's navigate to remotedesktop.google.com and click Access my computer. Login with your previously set up credentials. Once done, we'll see the main screen for our remote desktop menu. In the middle, you can see the computer we just set up for remote access. Let's give it a go and select it. After that, you'll be prompted for a PIN. You can choose to remember it on this device; however, if you're focused on security, you can choose not to remember it.
Great, it looks like our remote connection worked. We have access to the desktop and can do everything on it just like you'd be sitting at this desk, whether it's just surfing the internet or even launching different applications, like system information.
That's it for this demo. In this demo, we showed you how you can leverage a way to connect to your PC remotely without using a VPN.

5.5.7 Virtual Private Network Facts

This lesson covers the following topics:

- VPN basics
- VPN and wireless networks

VPN Basics

A virtual private network (VPN) is a remote-access connection that uses encryption to securely send data over an untrusted network. By using a VPN, you can take advantage of an existing internet connection to securely communicate between devices. When working with VPNs, consider the following:

- A VPN provides an alternative to:
  - WAN connections.
  - Connections that use telephone lines and a remote access server.
- VPNs work by using a Tunneling Protocol that encrypts packet contents and encapsulates those packets.
  - The encapsulated packets are routed through the internet using the information in the packet header.
  - When the packet reaches the destination device, the outer wrapping encapsulating the packets and the encryption is removed.
  - Only the destination device is allowed to remove the wrapping and restore the packet to its original form.
- The following are two styles of VPN tunnels commonly used:
  - Full tunnel, which routes all of a user's network traffic through the VPN tunnel. This can sometimes send traffic that is not necessary.
  - Split tunnel, which routes only certain types of traffic, usually determined by destination IP address, through the VPN tunnel. All other traffic is passed through the normal internet connection.
- VPNs can be implemented in the following ways:
  - A host-to-host VPN allows an individual host connected to the internet to establish a VPN connection to another host on the internet. Both devices must be configured for a VPN connection and have the software to encrypt and encapsulate the packets.
  - A site-to-site VPN uses routers on the edge of each site. The routers are configured for a VPN connection and encrypt and decrypt the packets being passed between the sites. With this configuration, individual hosts are unaware of the VPN.
  - A remote-access VPN uses a server (called a VPN concentrator) configured to accept VPN connections from individual hosts.
    - The VPN concentrator is located on the edge of a network.
    - The VPN concentrator establishes multiple connections with multiple hosts.
    - The individual hosts must be able to establish a VPN connection.
    - The hosts can access resources on the VPN server or the private network using the VPN connection.
  - An always-on VPN employs the concept that a user is always on the VPN, whether physically within the LAN or remotely. There is no turning it on or off. All traffic is basically fully tunneled.
- Tunnel endpoints are devices that can encrypt and decrypt packets. When you create a VPN, you establish a security association between the two tunnel endpoints. These endpoints create a secure virtual communication channel. Only the destination tunnel endpoint can unwrap packets and decrypt the packet contents.
- Routers use the decrypted packet headers to deliver the packet to the destination device. Intermediate routers along the path cannot read the encrypted packet contents.

When implementing a VPN, be sure to:

- Select a protocol that is supported by all devices that need to encrypt and encapsulate packets.
- Open the appropriate ports to allow VPN traffic through the firewall.

VPN and Wireless Networks

VPNs can also be used to help secure connections made over open wireless networks. Many establishments, such as airports, hotels, and restaurants, provide unsecured public Wi-Fi access. Because encryption is not used to secure the wireless connection, many users are hesitant to use these networks. In most cases, this hesitancy is warranted. However, it is generally considered acceptable to use a VPN connection to securely transfer data over an open Wi-Fi network. As long as strong tunneling ciphers and protocols are used, the VPN provides sufficient encryption to secure the connection even though the wireless network itself is not encrypted. It is recommended that you use IPsec or SSL to secure the VPN because these protocols are relatively secure. Avoid using PPTP with MS-CHAPv2 as this configuration setup is no longer considered secure.

If you are using a VPN over an open wireless network and need to access a secure website, be sure your browser's HTTPS requests go through the VPN connection. To conserve VPN bandwidth and improve latency, many VPN solutions automatically reroute web browsing traffic through the client's default network connection instead of through the VPN tunnel. This behavior would result in HTTP/HTTPS traffic being transmitted over the insecure open wireless network instead of through the secure VPN tunnel.

5.5.8 VPN Protocol Facts

This lesson covers the following topics:

- VPN Tunneling Protocol
- Transport Layer Security Tunneling
- Internet Protocol Security Tunneling
- Secure Communication and Access

VPN Tunneling Protocol

A VPN can use a Tunneling Protocol that encrypts packet contents and wraps them in an unencrypted packet. The Tunneling Protocol (also referred to as the VPN Protocol) identifies the methods that devices use to establish the VPN connection and encrypt the data. The three types of protocols used by VPNs are:

- Carrier Protocol (such as IP).
- Tunneling Protocol (such as PPTP or L2TP).
- Passenger Protocol (for the data being transmitted).

Many networks make use of a piece of hardware called a VPN concentrator. VPN concentrators are advanced routers that can create and maintain many secure connections to the network through VPN tunnels.

Transport Layer Security Tunneling

A transport layer security (TLS) VPN means the client connects to the remote access server using digital certificates. The server certificate identifies the VPN gateway to the client. Optionally, the client can also be configured with its own certificate. This allows for mutual authentication, where both server and client prove their identity to one another. TLS creates an encrypted tunnel for the user to submit authentication credentials. These would normally be processed by a RADIUS server. Once the user is authenticated and the connection is fully established, the VPN gateway tunnels all communications for the local network over the secure socket.

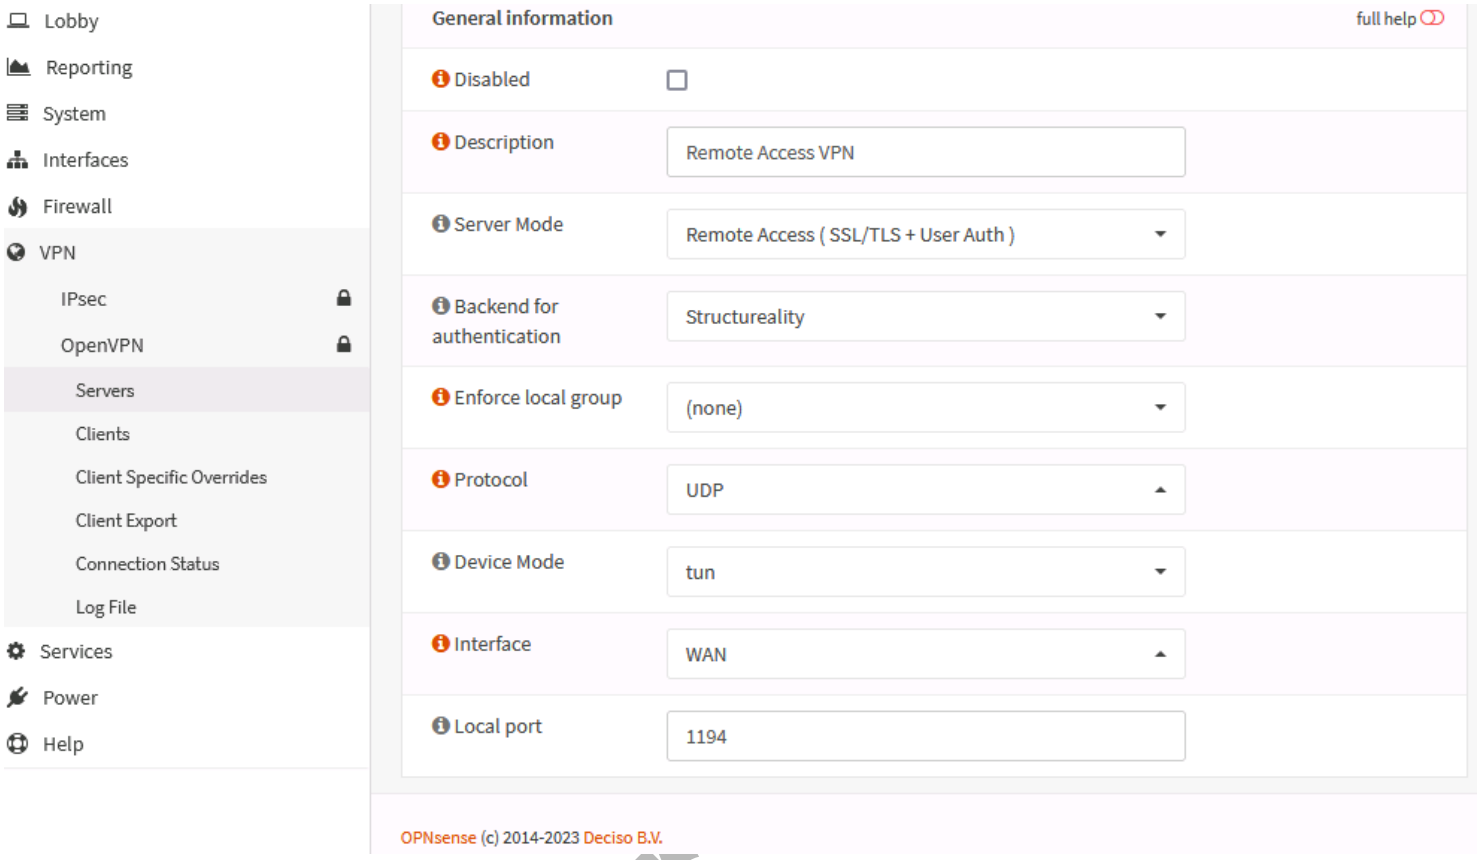

Configuring an OpenVPN server in the OPNsense security appliance. This configuration creates a remote access VPN. Users are authenticated via a RADIUS server on the local network. (Screenshot courtesy of OPNsense)

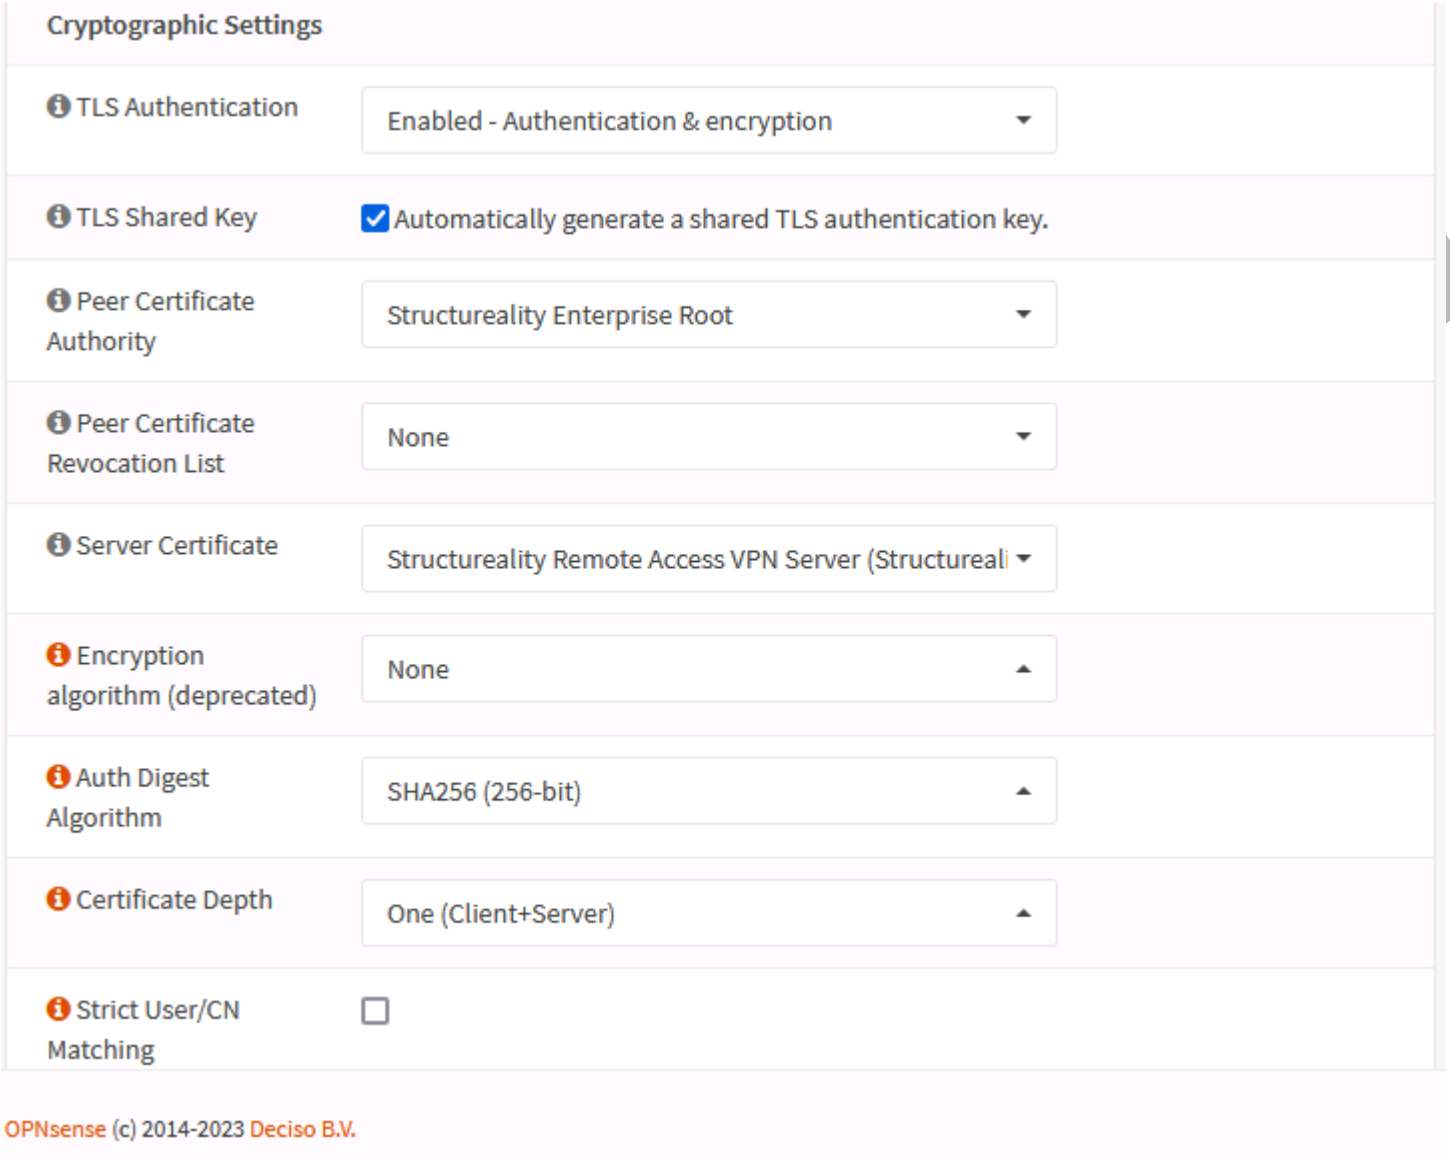

A TLS VPN can use either TCP or UDP. UDP might be chosen for marginally superior performance, especially when tunneling latency-sensitive traffic such as voice or video. TCP might be easier to use with a default firewall policy. TLS over UDP is also referred to as Datagram TLS (DTLS).

It is important to use a secure version of TLS. The latest version at the time of writing is TLS 1.3. TLS 1.2 is also still supported. Versions earlier than this are deprecated.

Internet Protocol Security Tunneling

Transport Layer Security is applied at the application level. Internet Protocol Security (IPsec) operates at the network layer of the OSI model (layer 3). This means that it can be implemented without having to configure specific application support and that it incurs less packet overhead.

There are two core protocols in IPsec, which can be applied singly or together, depending on the policy:

- Authentication Header (AH)—performs a cryptographic hash on the whole packet, including the IP header, plus a shared secret key (known only to the communicating hosts), and adds this value in its header as an

Integrity Check Value (ICV). The recipient performs the same function on the packet and key and should derive the same value to confirm that the packet has not been modified. The payload is not encrypted so this protocol does not provide confidentiality.

- Encapsulating Security Payload (ESP) can be used to encrypt the packet rather than simply calculating an ICV. ESP attaches three fields to the packet: a header, a trailer (providing padding for the cryptographic function), and an Integrity Check Value. Unlike AH, ESP excludes the IP header when calculating the ICV.

IPsec can be used in two modes:

- Transport mode—is used to secure communications between hosts on a private network. When ESP is applied in transport mode, the IP header for each packet is not encrypted, just the payload data. If AH is used in transport mode, it can provide integrity for the IP header.

- Tunnel mode—is used for communications between VPN sites across an unsecure network. With ESP, the whole IP packet (header and payload) is encrypted and encapsulated as a datagram with a new IP header. AH has no use case in tunnel mode, as confidentiality is usually required.

IPsec datagram using ESP in tunnel mode.

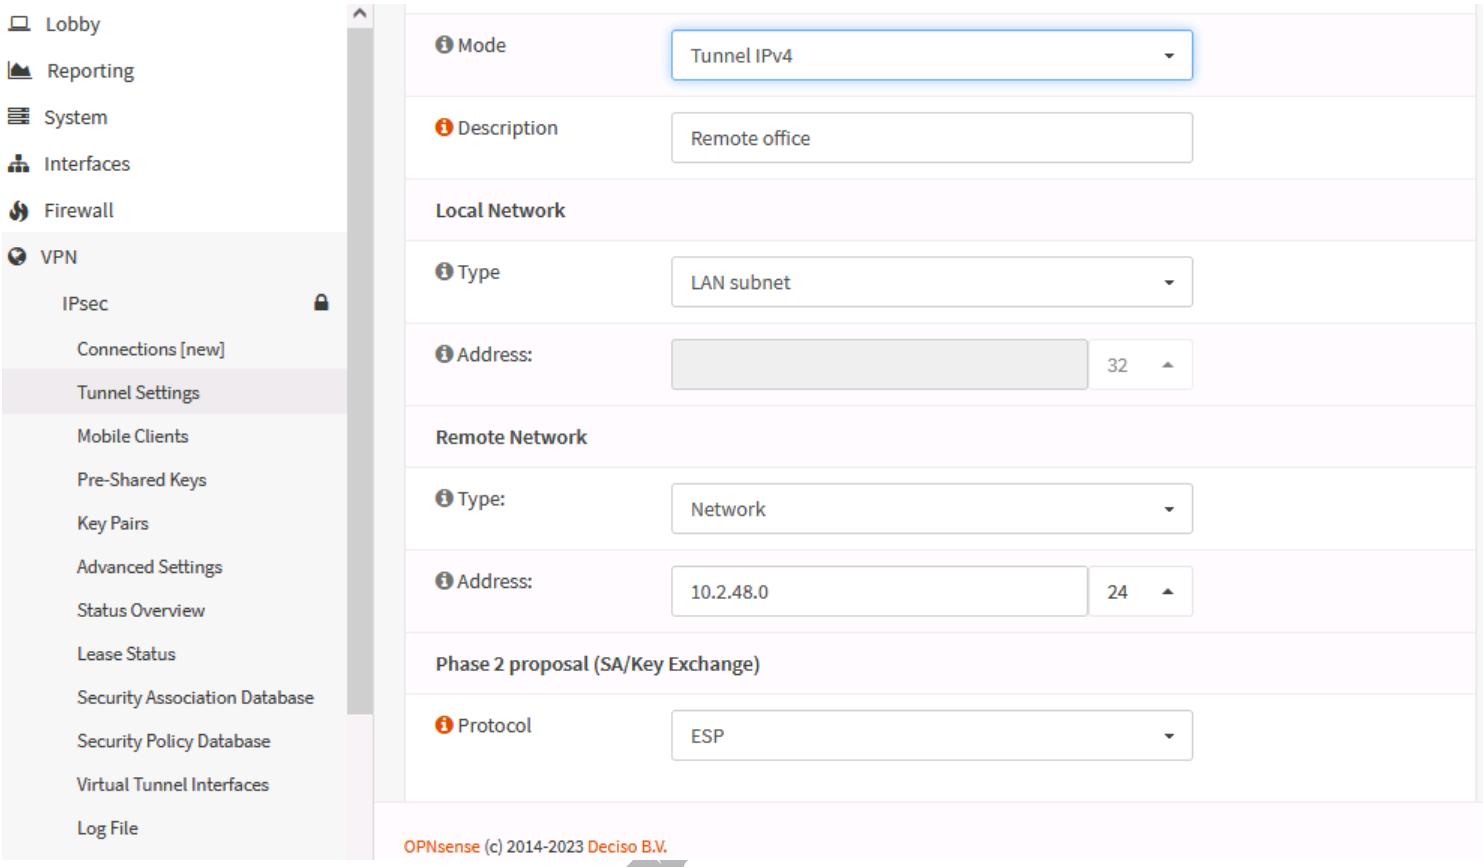

Configuring a site-to-site VPN using IPsec tunneling with ESP encryption in the OPNsense security appliance. (Screenshot courtesy of OPNsense)

Each host or router that uses IPsec must be assigned a policy. An IPsec policy sets the authentication mechanism and also the use of AH/ESP and transport or tunnel mode for a connection between two peers.

IPsec's encryption and hashing functions depend on a shared secret. The secret must be communicated to both peers, and the peers must perform mutual authentication to confirm one another's identity. The Internet Key Exchange (IKE) p rotocol implements an authentication method, selects which cryptographic ciphers are mutually supported by both peers, and performs key exchange. The set of properties is referred to as a s ecurity association (SA).

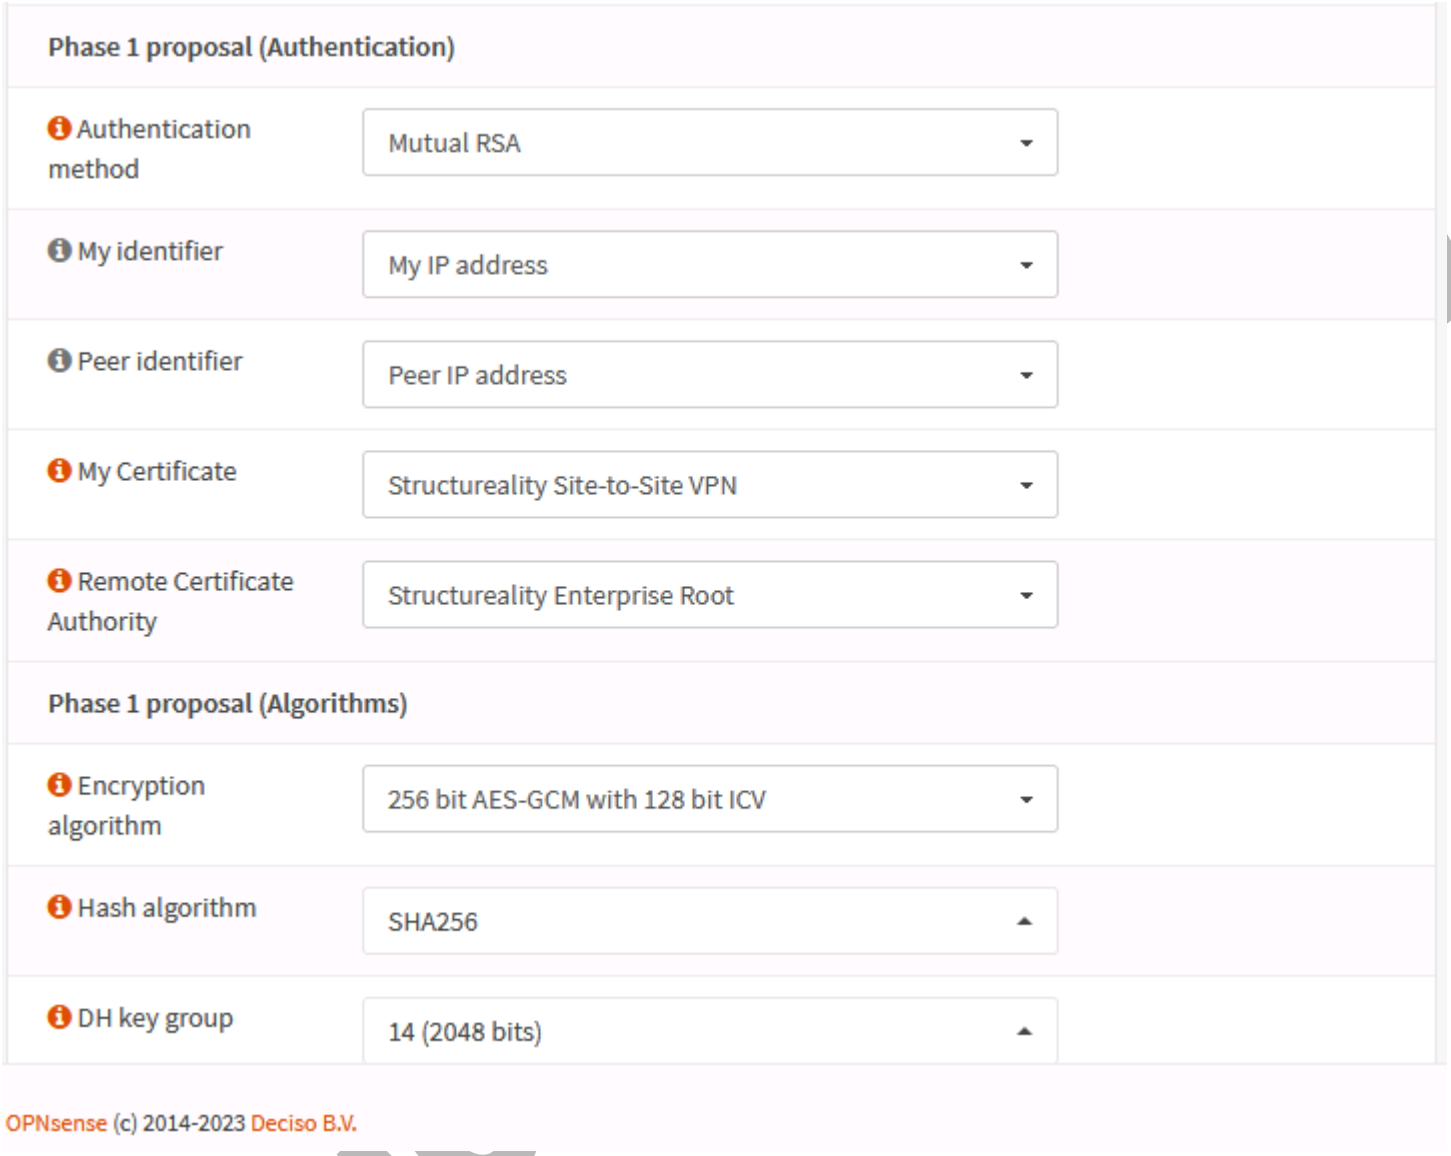

IKE negotiations take place over two phases:

- Phase I establishes the identity of the two peers and performs key agreement using the Diffie-Hellman algorithm to create a secure channel. Two methods of authenticating peers are commonly used:
- Digital certificates —are issued to each peer by a mutually trusted certificate authority to identify one another.
- Pre-shared key (group authentication) —is when the same passphrase is configured on both peers.
- Phase II uses the secure channel created in Phase I to establish which ciphers and key sizes will be used with AH and/or ESP in the IPsec session.

There are two versions of IKE. Version 1 was designed for site-to-site and host-to-host topologies and requires a supporting protocol to implement remote access VPNs. IKEv2 has some additional features that have made the protocol popular for use as a stand-alone remote access client-to-site VPN solution. The main changes are the following:

- Supports EAP authentication methods, allowing, for example, user authentication against a RADIUS server.
- Provides a simple setup mode that reduces bandwidth without compromising security.

- Allows network address translation (NAT) traversal and MOBIKE multihoming. NAT traversal makes it easier to configure a tunnel allowed by a home router/firewall. Multihoming means that a smartphone client with Wi-Fi and cellular interfaces can keep the IPsec connection alive when switching between them.

Secure Communication and Access

A Software-Defined Wide Area Network (SD-WAN) enables organizations to connect their various branch offices, datacenters, and cloud infrastructure over a wide area network (WAN). One of the key advantages of SD-WAN is its ability to provide enhanced security features and considerations. For example, SD-WAN uses encryption to protect data as it travels across the network and can segment network traffic based on priority ratings to ensure that critical data is fully protected.

Additionally, SD-WAN can intelligently route traffic based on the application and tightly integrate with firewalls to provide additional protection against known threats. SD-WAN centralizes the management of network security policies to simplify enforcing security measures across an entire network.

Secure Access Service Edge (SASE) combines the protection of a secure access platform with the agility of a cloud-delivered security architecture. SASE offers a centralized approach to security and access, providing end-to-end protection and streamlining the process of granting secure access to all users, regardless of location. SASE is a network architecture that combines wide area networking (WAN) technologies and cloud-based security services to provide secure access to cloud-based applications and services.

SASE offers several security features to help organizations protect their networks and data as SASE operates under a zero trust security model. SASE incorporates Identity and Access Management (IAM) and assumes all users and devices are untrusted until they are authenticated and authorized. SASE also provides a range of threat prevention features, such as intrusion prevention, malware protection, and content filtering.

5.5.9 Implement Secure Remote Access Protocols

5.5.10 Practice Questions (Section Quiz)

q_vpn_config_secp8

A salesperson in your organization spends most of her time traveling between customer sites. After a customer visit, she must complete various managerial tasks, such as updating your organization's order database.

Because she rarely comes back to your home office, she usually accesses the network from her notebook computer using Wi-Fi access provided by hotels, restaurants, and airports.

Many of these locations provide unencrypted public Wi-Fi access, and you are concerned that sensitive data could be exposed. To remedy this situation, you decide to configure her notebook to use a VPN when accessing the home network over an open wireless connection.

Which key steps should you take when implementing this configuration? (Select two.)

Answers:

- ***Configure the VPN connection to use IPsec**
- Configure the VPN connection to use MS-CHAPv2
- ***Configure the browser to send HTTPS requests through the VPN connection**
- Configure the browser to send HTTPS requests directly to the Wi-Fi network without going through the VPN connection
- Configure the VPN connection to use PPTP

Explanation:

It is generally considered acceptable to use a VPN connection to securely transfer data over an open Wi-Fi network. As long as strong tunneling ciphers and protocols are used, the VPN provides sufficient encryption to secure the connection, even though the wireless network itself is not encrypted. It is recommended that you use IPsec or SSL to secure the VPN, as these protocols are relatively secure. You should also configure the browser's HTTPS requests to go through the VPN connection. To conserve VPN bandwidth and improve latency, many VPN solutions automatically reroute web browsing traffic through the client's default network connection instead of through the VPN tunnel. This behavior would result in HTTP/HTTPS traffic being transmitted over the unsecure open wireless network instead of though the secure VPN tunnel.

Avoid using PPTP with MS-CHAPv2 in a VPN over open wireless configuration, as these protocols are no longer considered secure.

q_vpn_example_secp8

An IT team at a global pharmaceutical company has decided to implement a virtual private network (VPN) for remote employees to securely access internal company resources from home.

Which of the following is a primary reason for this decision?

Answers:

- ***VPNs encrypt and encapsulate all traffic to create a secure tunnel**
- VPNs provide a direct connection to a single machine.
- VPNs provide secure command-line access and file transfer.
- VPNs are designed for managing devices on IP networks.

Explanation:

VPNs create a secure private connection between the remote user's device and the company's internal network, require authentication, and use a network tunnel to encrypt and encapsulate all traffic.

Remote Desktop Protocol (RDP) allows a user to remotely access and control another machine but does not encrypt all traffic or create a network tunnel.

This rationale is characteristic of Secure Shell (SSH), which offers secure, encrypted command-line access and file transfer but does not encapsulate all traffic or create a network tunnel.

Simple Network Management Protocol (SNMP) is a device management protocol for IP networks, not for creating a secure connection for a remote user.

q_vpn_remote_secp8

A group of salesmen would like to remotely access your private network through the internet while they are traveling. You want to control access to the private network through a single server.

Which solution should you implement?

Answers:

- ***VPN concentrator**
- IDS

- IPS
- DMZ

Explanation:

With a remote access VPN, a server on the edge of a network (called a VPN concentrator) is configured to accept VPN connections from individual hosts. Hosts that are allowed to connect using the VPN connection are granted access to resources on the VPN server or the private network.

A demilitarized zone (DMZ), also called a screened subnet, is a buffer network (or subnet) that sits between the private network and an untrusted network (such as the internet). A RADIUS server is used to centralize authentication, authorization, and accounting for multiple remote access servers. However, clients still connect to individual remote access servers.

An intrusion detection system (IDS) is a special network device that can detect attacks and suspicious activity. A passive IDS monitors, logs, and detects security breaches, but it does not take action to stop or prevent an attack. An active IDS (also called an intrusion protection system or IPS) performs the functions of an IDS but can also react when security breaches occur.

q_vpn_secure_secp8

A VPN is primarily used for which of the following purposes?

Answers:

- Allow remote systems to save on long-distance charges
- ***Support secured communications over an untrusted network**
- Allow the use of network-attached printers
- Support the distribution of public web documents

Explanation:

A VPN (virtual private network) is used primarily to support secured communications over an untrusted network. A VPN can be used over a local area network, across a WAN connection, over the internet, and even between a client and a server over a dial-up internet connection.

All of the other items listed in this question are benefits or capabilities that are secondary to this primary purpose.

q_vpn_site_secp8

Which VPN implementation uses routers on the edge of each site?

Answers:

- ***Site-to-site VPN**
- Host-to-host VPN
- Remote access VPN
- Always-on VPN

Explanation:

A site-to-site VPN uses routers on the edge of each site. The routers are configured for a VPN connection and encrypt and decrypt the packets being passed between the sites. With this configuration, individual hosts are unaware of the VPN.

A host-to-host VPN allows an individual host connected to the internet to establish a VPN connection to another host on the internet. Both devices must be configured for a VPN connection and have the software to encrypt and encapsulate the packets.

A remote access VPN uses a server (called a VPN concentrator) configured to accept VPN connections from individual hosts.

An always-on VPN employs the concept that a user is always on the VPN, whether physically within the LAN or remotely. There is no turning it on or off. All traffic is basically fully tunneled.

q_vpn_split_secp8

Which VPN tunnel style routes only certain types of traffic?

Answers:

- Full
- ***Split**
- Site-to-site
- Host-to-host

Explanation:

A VPN split tunnel routes only certain types of traffic, usually determined by destination IP address, through the VPN tunnel. All other traffic is passed through the normal internet connection.

A full VPN tunnel routes all of a user's network traffic through the VPN tunnel. This can sometimes send traffic that is not necessary.

A site-to-site VPN is a VPN implementation that uses routers on the edge of each site.

A host-to-host VPN implementation allows an individual host connected to the internet to establish a VPN connection to another host on the internet.

q_vpn_prot_esp_01_secp8

Which IPSec subprotocol provides data encryption?

Answers:

- SSL
- AES
- AH
- ***ESP**

Explanation:

Encapsulating Security Payload (ESP) Protocol provides data encryption for IPSec traffic.

Authentication Header (AH) provides message integrity through authentication, verifying that data is received unaltered from the trusted destination. AH provides no privacy and is often combined with ESP to achieve integrity and confidentiality.

SSL is standard technology for securing an internet connection by encrypting data sent between a website and a browser. It is not used for IPSec.

The Advanced Encryption Standard (AES) is an algorithm that uses the same key to encrypt and decrypt protected data. However, it is not used by IPSec.

q_vpn_prot_esp_02_secp8

In addition to Authentication Header (AH), IPSec is comprised of what other service?

Answers:

- Encryption File System (EFS)
- Extended Authentication Protocol (EAP)
- Advanced Encryption Standard (AES)
- ***Encapsulating Security Payload (ESP)**

Explanation:

IPsec is comprised of two services. One service is named Authentication Header (AH), and the other named Encapsulating Security Payload (ESP). AH is used primarily for authenticating the two communication partners of an IPsec link. ESP is used primarily to encrypt and secure the data transferred between IPsec partners. IPSec employs ISAKMP for encryption key management.

EFS provides a file system interface, file system access semantics (such as strong consistency and file locking), and concurrently accessible storage for up to thousands of EC2 instances. It is not used with IPSec.

Extensible Authentication Protocol (EAP) is used to pass the authentication information and handles and defines an authentication. It is not used with IPSec.

The Advanced Encryption Standard (AES) is an algorithm that uses the same key to encrypt and decrypt protected data. However, it is not used with IPSec.

q_vpn_prot_ike_ipsec_secp8

A network engineer has the task of creating a remote access solution for a global enterprise. The solution should secure encrypted communication for the company's employees worldwide and detect potential security threats in real time.

Which configuration should the network engineer deploy to meet these requirements?

Answers:

- ***A VPN utilizing IKE and IPSec protocols, combined with an inline intrusion detection system (IDS)**
- A network fortified by 802.1X port security, an Extensible Authentication Protocol (EAP), and a load balancer
- A network equipped with a Next Generation Firewall (NGFW), a Web Application Firewall (WAF), and an intrusion prevention system (IPS) in tap/monitor mode
- A Software-Defined Wide Area Network (SD-WAN) with secure access service edge (SASE) implementation, supplemented by an intrusion prevention system (IPS)

Explanation:

A VPN using internet key exchange and IPSec protocols secures remote access and communication. An inline IDS detects potential security threats in real time.

802.1X port security and EAP provide robust network security but don't address secure remote access. A load balancer manages traffic but doesn't contribute to threat detection.

An NGFW and WAF protect the network and web applications but don't provide secure remote access. An IPS in tap/monitor mode detects intrusions but isn't as effective as an inline system.

SD-WAN with SASE provides secure, scalable cloud-based network architecture but focuses more on optimization than secure remote access. An IPS prevents intrusions but doesn't directly address threat detection or secure remote access.

q_vpn_prot_ike_phases_secp8

Which of the following is commonly used in the first phase of Internet Key Exchange (IKE) negotiations for authenticating the identity of peers?

Answers:

- ***Digital certificates**
- Passwords
- Biometrics
- Security questions

Explanation:

Digital certificates is the correct answer. Digital certificates are commonly used in the first phase of IKE negotiations to authenticate the identity of peers. They provide a way to exchange public keys and contain information about the entity it represents, the entity that issued the certificate, and the digital signature of the issuer.

While passwords can be used in some forms of authentication, they are not typically used in the first phase of IKE negotiations. The use of passwords alone would not provide the level of security required for VPN connections.

Biometric authentication, such as fingerprint or facial recognition, is not used in IKE negotiations. This form of authentication is more commonly used for physical or device access rather than network connections.

Security questions are typically used for resetting passwords or additional layer of authentication, not for the initial authentication of peers in IKE negotiations.

q_vpn_prot_ipsec_01_secp8

Which statement BEST describes IPsec when used in tunnel mode?

Answers:

- ***The entire data packet, including headers, is encapsulated**
- The identities of the communicating parties are not protected
- IPsec in tunnel mode may not be used for WAN traffic
- Packets are routed using the original headers, and only the payload is encrypted

Explanation:

When using IPsec in tunnel mode, the entire data packet, including original headers, is encapsulated. New encrypted packets are created with headers indicating only the endpoint addresses. Tunneling protects the identities of the communicating parties and original packet contents.

Tunneling is frequently used to secure traffic traveling across insecure public channels, such as the internet. IPsec in tunnel mode is the most common configuration for gateway-to-gateway communications.

In transport mode, routing is performed using the original headers; only the packet's payload is encrypted. Transport mode is primarily used in direct host-to-host communication outside of a dedicated IPsec gateway/firewall configuration.

q_vpn_prot_ipsec_02_secp8

A security team in a multinational organization decides to improve the security of their inter-office communications. They agree to use a tunneling protocol that can offer confidentiality, sender authentication, and message integrity.

They need a protocol that operates at the network level.

Which protocol BEST fulfills the team's requirements for securing inter-office communications and operates at the network level?

Answers:

- HTTPS
- SSH
- ***IPSec**
- TLS

Explanation:

Internet Protocol Security (IPSec) best meets the team's requirements by providing confidentiality, sender authentication, and message integrity by functioning at the network level.

Hypertext Transfer Protocol Secure (HTTPS) operates at the application layer, primarily securing communication between web browsers and web servers, but it does not cater to inter-office communications.

Secure Shell (SSH) is a network communication protocol that allows two computers or devices to communicate and share data. However, it is not normally associated with securing inter-office communications.

Transport Layer Security (TLS) is a security protocol designed to facilitate privacy and data security for communications over the Internet, not for inter-office communications.

q_vpn_prot_l2tp_secp8

Which VPN protocol typically employs IPsec as its data encryption mechanism?

Answers:

- PPTP
- ***L2TP**
- PPP
- L2F

Explanation:

L2TP (Layer 2 Tunneling Protocol) is the VPN protocol that typically employs IPsec as its data encryption mechanism. L2TP is the recommended VPN protocol to use on dial-up VPN connections.

PPTP and PPP only support CHAP and PAP for data encryption. L2F offers no data encryption.

q_vpn_prot_sd-wan_secp8

A global enterprise is expanding its infrastructure to handle the increased service demand. The enterprise plans to deploy several load balancers across its geographically distributed data centers as part of the expansion process.

The senior network engineer's responsibilities include designing a secure load-balancing solution that ensures maximum availability, minimizes the attack surface, and provides secure communication between the data centers.

What approach BEST ensures secure and efficient load balancing across geographically distributed data centers?

Answers:

- ***Use a Software-Defined Wide Area Network (SD-WAN) solution along with the load balancers.**
- Deploy load balancers in each data center and secure them using a Layer 7 firewall.
- Implement secure access service edge (SASE) to provide secure and enhanced network connectivity.
- Configure load balancers in fail-closed mode and use a VPN for secure communication.

Explanation:

SD-WAN provides secure, efficient, and flexible load balancing across geographically distributed data centers.

Deploying load balancers in each data center and securing them with a Layer 7 firewall may not be the best solution for secure and efficient load balancing.

SASE provides secure network connectivity but does not directly address the load-balancing requirements of geographically distributed data centers.

Configuring load balancers in a fail-closed mode mitigates the risk of unfiltered traffic but might interrupt service availability. Using a VPN may not be the most efficient solution due to potential latency issues.

q_vpn_prot_tls_01_secp8

The IT department in a large multinational corporation faces challenges managing secure communications for remote desktop connections. The increasing number of remote employees has made it essential to ensure that their remote desktop connections are secure. The IT department is considering various measures to establish secure communication.

Given the challenges the corporation faces, what approach should the IT department adopt to ensure secure communications for remote desktop connections while maintaining the manageability and performance of the enterprise infrastructure?

Answers:

- ***Implement TLS for all remote desktop connections**
- Disable all firewall rules for remote desktop connections
- Establish VPN tunnels for all users without using any encryption protocols
- Enable open access to all remote desktop connections for easy manageability

Explanation:

Transport layer security (TLS) provides secure communication for remote desktop connections by encrypting the data transmitted between the end user and the remote desktop server, reducing the risk of data breaches.

Disabling all firewall rules for remote desktop connections is not a secure practice. It exposes the enterprise infrastructure to potential cyberattacks.

Establishing virtual private network (VPN) tunnels for all users without using any encryption protocols does not secure communications. While VPN tunnels provide a layer of security, the data transmitted is still vulnerable to interception and breaches without encryption.

Enabling open access to all remote desktop connections for easy manageability compromises the security of the enterprise infrastructure. It exposes the enterprise to numerous potential security threats and cyberattacks.

q_vpn_prot_tls_02_secp8

A medium-sized organization implements a secure access service edge (SASE) solution to secure its distributed network. The IT manager wants to protect all data in transit.

Which of the following should the organization implement alongside SASE to achieve this goal?

Answers:

- ***A Transport Layer Security (TLS) protocol to secure data in transit**
- A standalone intrusion detection system (IDS) to monitor network traffic
- A Software-Defined Wide Area Network (SD-WAN) to manage network connectivity
- A Layer 4 firewall to control access to network resources

Explanation:

The Transport Layer Security (TLS) protocol encrypts data in transit, providing an additional layer of security alongside SASE to help to ensure sensitive information remains confidential and protected from unauthorized access.

While an intrusion detection system (IDS) can monitor network traffic for signs of security threats, it doesn't specifically secure data in transit, meaning data could still be vulnerable to interception or tampering.

An SD-WAN can manage network connectivity and optimize traffic routing but does not inherently secure data in transit, so additional measures may be necessary to protect sensitive information.

A Layer 4 firewall controls access to network resources by filtering traffic based on port numbers and protocols, but it does not specifically secure data in transit.

5.6 Network Access Control

As you study this section, answer the following questions:

- How do remediation servers and auto-remediation help clients become compliant?
- What are the security standards NAC uses for evaluation?
- What is an NAC agent? What types of NAC agents are available?
- What are the four steps of the NAC process?

The key terms for this section include:

Term	Definition
Network access control (NAC)	A general term for the collected protocols, policies, and hardware that authenticate and authorize access to a network at the device level.
Bring your own device (BYOD)	Security framework and tools to facilitate use of personally owned devices to access corporate networks and data.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.5 Given a scenario, modify enterprise capabilities to enhance security. • Network access control (NAC)

5.6.1 Network Access Control (Lesson Video)

Transcript:

Network Access Control, or NAC, is a process control that prevents unauthorized access and preserves the integrity of a production network. Depending on the configuration, an endpoint must meet certain criteria in order to join the internal network. In this lesson, we'll discuss the components and steps required for an end device to access an internal network.

Before an end device can connect, it must go through several validation checks, and it meet specified criteria. For example, a system may need to be running Windows 10 with update 1907. It may also need to have the latest antivirus latest definitions installed, have Windows firewall enabled, have the latest updates installed, and have automatic updates enabled.
If any of this criteria isn't met, the system belongs on a restricted network. A restricted network has equipment with the software, drivers, and updates you need to retrieve the required settings. Once the system is remediated and receives all the updates and system changes, it's safe for it to join the production network.
Depending on the company's business market, a compliance model may dictate the level of protection and integrity a network system requires. Compliance industries include the medical, defense, and financial sectors. Each one has government-regulated compliance rules.
The basic tenants of NAC include zero-day attack mitigation, role-based controls, traffic encryption, identity management, and policy enforcement. These goals vary widely depending on NAC implementation, configuration, compliance requirements, and industry. To create an effective solution, you have to understand the resources you're protecting and why you're protecting them.
Now, network access control is a process, not a product, although several companies implement a product that uses the NAC process. For example, Cisco has a solution called Identity Services Engine, or ISE, while Microsoft has Network Access Protection, or NAP. Each is a framework that provides protection against unauthorized access or rogue system access to a protected network. The distinction between products is vendor-specific, and so is the implementation.

Before a device can connect to the production network, it must go through a health check. The health check is configured by system management and may include criteria like antivirus software, firewall configurations, and patches. This health check is performed by software designed solely for this purpose.
One way to run this software is using an agent. An agent is a preinstalled software program that performs these health checks at predefined times. One example of how this might work happens when the workstation attempts a connection to the production network. At that time, the agent is activated and goes through its checks to ensure that the workstation meets the minimum requirements to access the production network.
Depending on the vendor, it's also possible to perform health checks on demand. This process is agentless, meaning it doesn't require a permanent agent to be preinstalled. Instead, the required software is downloaded and executed at the workstation when it attempts to connect to the network. Access to the production network is only granted if all the prerequisites are met and the health check passes.
Network Access Control is a policy-driven construct that maintains network integrity using authentication and authorization policies. NAC can be compartmentalized for different types of devices such as Internet of Things, Bring Your Own Device, or vendors and contractors.
NAC defines the prerequisites and identity, or authentication. If a devices meets requirements, NAC provides access to the target information. You can think of NAC as a layered approach to providing network admission and access. Let's see how it works.
The first layer defines the tasks required for authentication. The authentication layer defines all the prerequisites a device must meet to access the network. This criteria is detailed in policies that define minimum requirements for the device, such as anti-malware software, OS type, patch level, and so on.
Keep in mind that the policies may be different for dissimilar types of devices. For instance, a tablet may be required to be managed by an MDM package or have backup and remote wipe enabled. Likewise, an Internet of Things device might be required to use a secure VLAN.
If a device fails authentication, it's forwarded to a remediation network. This network is protected and doesn't have access to any production information or outside areas such as the internet. Instead, the remediation network is preconfigured to contain all the necessary software and procedures to bring a device into compliance. It may have software updates, anti-malware definitions, and other control software. Once a device goes through remediation, it goes through the authentication process again. If it passes, it moves on to the authorization process.
The most secure method for ensuring integrity is a zero-trust access model. This is the concept of least privilege, or implicit deny all. With this model, users or devices are only given the permissions they require to do their jobsâ€“ no more, no less.
The authorization process looks at the authentication information and applies the policies you choose. After the device goes through the authorization process, it's granted access to the production network. This process isn't easy to set up, and it's often set up incorrectly. Several policies must be put into place, and boundary networks must be configured. Everything from connectivity devices to routing and switching must comply with the necessary logic to ensure that this process works. While difficult, it may be absolutely necessary in higher security environments where integrity is required. This is especially true in regulated industries.
NAC is a complicated process that requires forethought. The first step is planning. A committee should convene and make decisions that define how NAC will work. Next, roles, identities, and permissions must be defined. Then these polices must be applied. As business needs change, NAC configuration must be reviewed to determine whether it should change, too.
That's it for this lesson. In this video, we discussed how Network Access Control is a concept that helps to ensure the integrity of a network system. It provides a more granular approach to providing network access by enforcing minimum requirements applied to devices.
You can define items such as anti-malware, OS versions, patch levels, and more. If devices don't meet these requirements, they're redirected to a remediation network. Once they're authenticated, devices are authorized and granted access. This is an ongoing process as business processes change and the organization grows.

5.6.2 Network Access Control Facts

This lesson covers the following topics:

- NAC overview
- Agent vs. agentless configurations
- NAC process

NAC Overview

Network access control (NAC) not only authenticates users and devices before allowing them access to the network but also checks and enforces compliance with established security policies. NAC ensures that devices meet a minimum set of security standards before being granted network access by evaluating:

- The operating system version
- Patch level
- Antivirus status
- The presence of specific security software

To ensure users and devices can only access the resources necessary to complete their duties, NAC can also restrict access based on:

- User profile
- Device type
- Location
- Other attributes

NAC plays a crucial role in identifying and quarantining suspicious or non-compliant devices. For organizations with bring-your-own-device (BYOD) policies and increasing use of IoT devices, NAC helps organizations secure their internal network environment against unauthorized access.

NAC and virtual local area networks (VLANs) work together to improve and automate network security. One of the primary ways NAC integrates with VLAN protections is through dynamic VLAN assignment. Dynamic VLAN assignment is a NAC feature that assigns a VLAN to a device based on the user's identity attributes, device type, device location, or health check results. For instance, a visiting user (such as a vendor) might be placed into a VLAN that only provides internet access, while a corporate user would be assigned to a VLAN with access to internal resources. Additionally, NAC can interact with dynamic VLAN to implement quarantine procedures. If a device is non-compliant with security policies—for example if it lacks updated antivirus software—the NAC system can automatically move it to a quarantine VLAN. This VLAN is generally isolated from the rest of the network, limiting potential damage from threats like malware.

NAC is often integrated with automatic remediation systems, which helps bring a computer into compliance when NAC discovers missing requirements. When such a system is used, a connecting device only has access to resources that can resolve or remediate the issues. After the issues have been remediated, the device is reevaluated and then allowed to access network resources, like the internet or databases.

The NAC process is usually accomplished using a two-stage process of authentication and authorization. If the requirements for either of these stages are not met, the access request is denied. This is often referred to as zero-trust security, meaning nothing is trusted unless it can pass both the authentication and authorization stages.

- Authentication defines all the prerequisites a device must meet to access the network. This criteria is detailed for such things as anti-malware, OS, patch level, and so on.
- Authorization looks at the authentication information and applies the appropriate policies to provide the device with the access it’s defined to receive.

Agent vs. Agentless Configurations

NAC can enforce security policies using agent-based and agentless methods. In an agent-based approach, a software agent is installed on the devices that connect to the network. This agent communicates with the NAC platform, providing detailed information about the device's status and compliance level. An agent-based NAC implementation can enable features such as automatic remediation, where the NAC agent can perform actions like updating software or disabling specific settings to bring a device into compliance with mandatory security configurations.

In contrast, an agentless NAC approach uses port-based network access control or network scans to evaluate devices. For example, agentless NAC may use DHCP fingerprinting to identify the type and configuration of a device when it connects, or it might perform a network scan to detect open ports or active services. While agentless methods may not provide as detailed information about a device's status, they can be used with any device that connects to the network, including guest or IoT devices, without requiring any prior configuration.

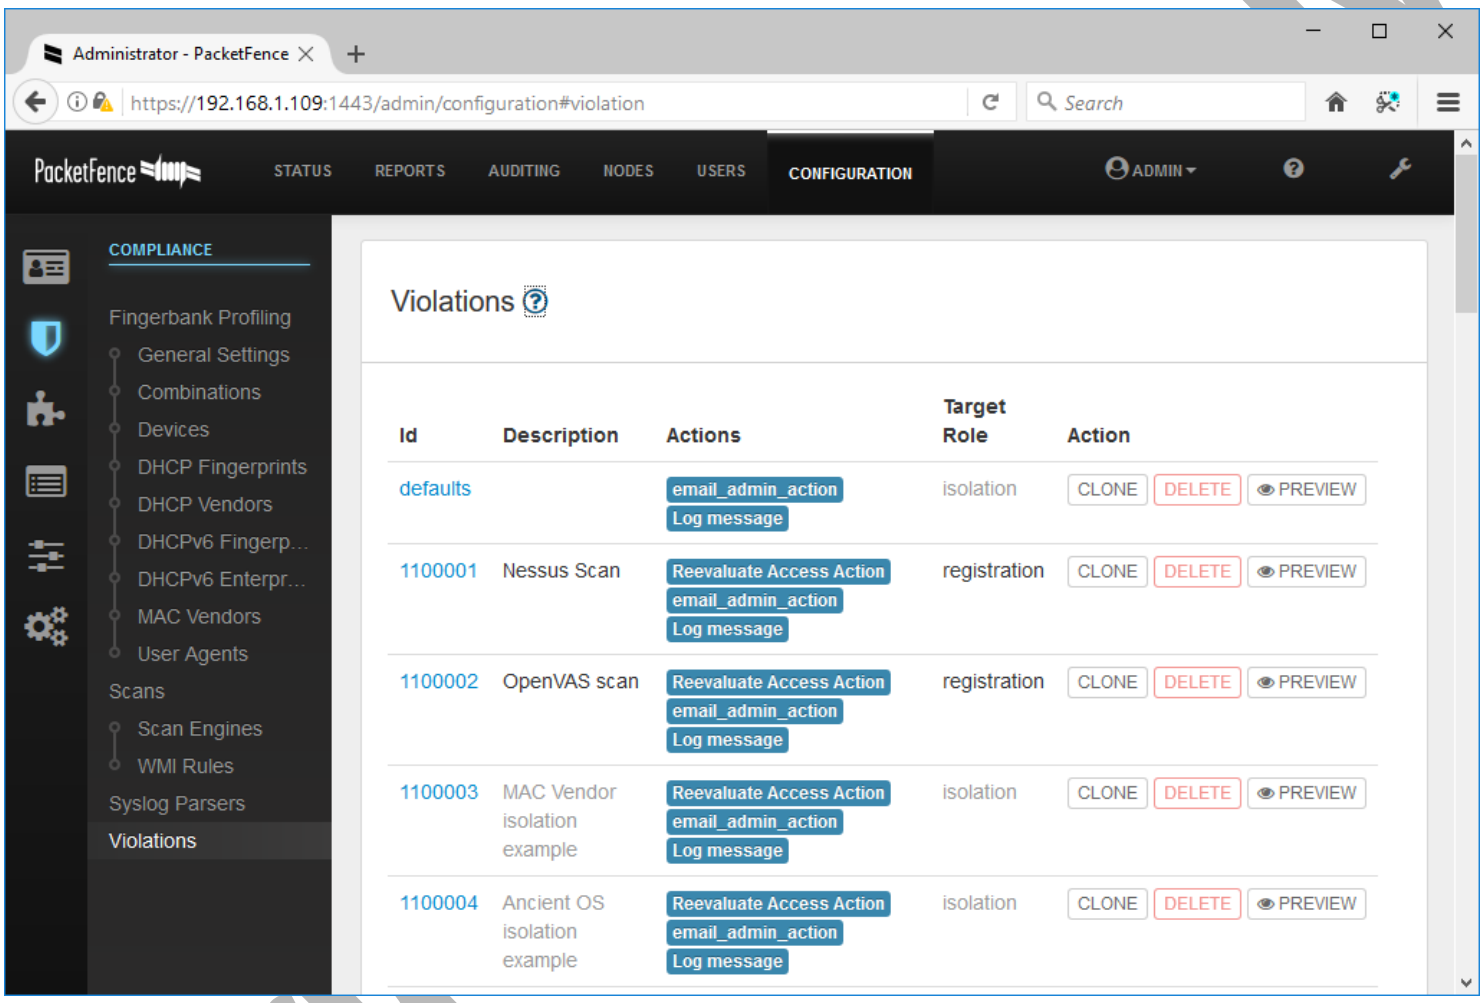

Defining policy violations in PacketFence Open Source NAC. (Screenshot used with permission from packetfence.org)

An agent can be persistent, in which case it is installed as a software application on the client or nonpersistent. A nonpersistent (or dissolvable) agent is loaded into memory during posture assessment but is not installed on the device.

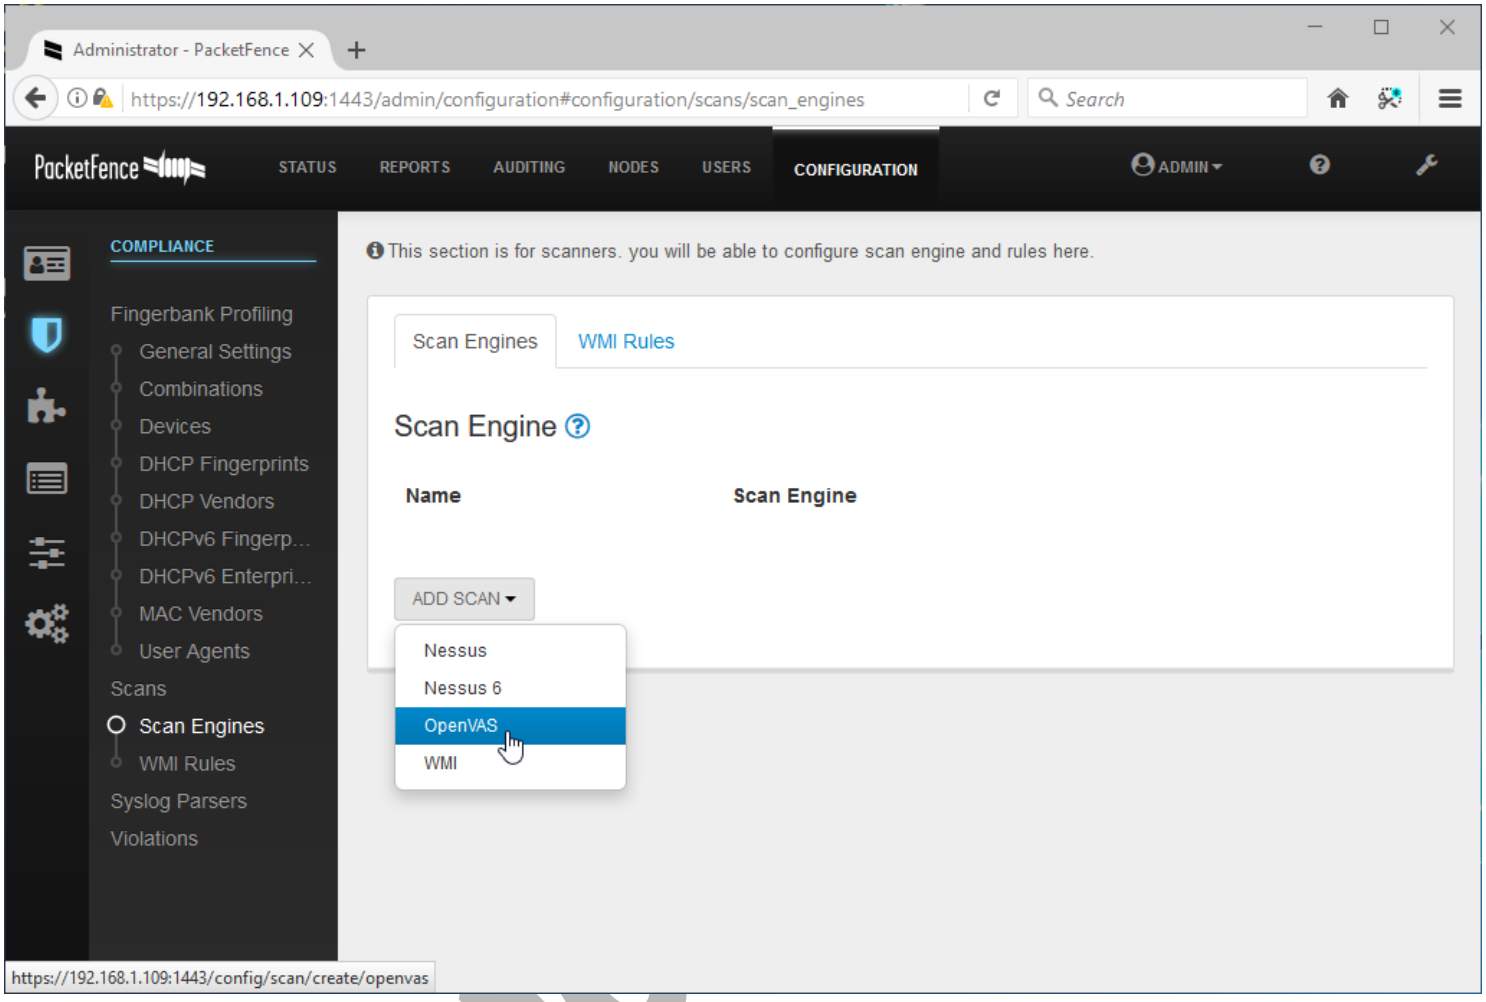

PacketFence supports the use of several scanning techniques, including vulnerability scanners, such as Nessus and OpenVAS, Windows Management Instrumentation (WMI) queries, and log parsers. (Screenshot used with permission from packetfence.org)

NAC Process

NAC is a complicated process that requires forethought. Use the following process when implementing NAC:

- Plan - A committee should convene and make decisions that define how NAC should work.
- Define - The roles, identities, and permissions (policies) must be defined.
- Apply - Once defined, the policies must be applied.
- Review/Revise - As business needs change, the process must be reviewed to determine whether changes are required.

5.6.3 Practice Questions (Section Quiz)

q_nac_agentless_secp8

Which of the following NAC agent types would be used for IoT devices?

Answers:

- Dissolvable
- ***Agentless**
- Zero-trust
- Permanent

Explanation:

An agentless agent is on the domain controller. When the user logs into the domain, it authenticates with the network. Agentless NAC is often used when there is limited disk space, such as for Internet of Things (IoT) devices.

A dissolvable agent is downloaded, or a temporary connection is established. The agent is removed once the user is done with it.

Zero-trust security means nothing is trusted unless it can pass both the authentication and authorization stages.

A permanent agent resides on a device permanently.

q_nac_apply_secp8

Which of the steps in the Network Access Control (NAC) implementation process occurs once the policies have been defined?

Answers:

- Plan
- ***Apply**
- Review
- Test

Explanation:

The third step in implementing NAC is to apply the policies. This occurs after the policies have been defined.

Planning is the first step in the NAC implementation process and needs to be done before defining the policies.

Review is the final step in the NAC implementation process. As business needs change, the process must be reviewed to determine whether changes are required.

Testing is not a step in the NAC implementation process.

q_nac_authentication_secp8

Which of the following defines all the prerequisites a device must meet in order to access a network?

Answers:

- Authorization
- ***Authentication**
- Zero-trust security
- Identity Services Engine (ISE)

Explanation:

Authentication defines all the prerequisites a device must meet in order to access a network. These criteria are detailed for such things as anti-malware, OS, and patch level.

Authorization looks at the authentication information and applies the appropriate policies to provide the device with the access it's defined to receive.

Zero-trust security means nothing is trusted unless it can pass both the authentication and authorization stages.

Identity Services Engine (ISE) is Cisco's NAC solution.

q_nac_authorization_secp8

Which of the following applies the appropriate policies in order to provide a device with the access it's defined to receive?

Answers:

- ***Authorization**
- Authentication
- Zero-trust security
- Identity Services Engine

Explanation:

Authorization looks at the authentication information and applies the appropriate policies in order to provide a device with the access it's defined to receive.

Authentication defines all the prerequisites a device must meet in order to access a network. These criteria are detailed for such things as anti-malware, OS, patch level, and so on.

Zero-trust security means nothing is trusted unless it can pass both the authentication and authorization stages.

Identity Services Engine (ISE) is Cisco's NAC solution.

q_nac_automatic_remediation_secp8

As a network administrator, you have implemented a Network Access Control (NAC) system with automatic remediation capabilities in your organization.

One day, you notice that a significant number of devices are being quarantined frequently by the NAC system due to non-compliance with security policies.

What should be your next course of action?

Answers:

- Ignore the issue as the NAC system is doing its job by quarantining non-compliant devices.
- Disable the automatic remediation feature to reduce the number of quarantined devices.
- ***Investigate the root cause of the frequent non-compliance and address it.**
- Reduce the security standards set by the NAC system to decrease the number of quarantined devices.

Explanation:

Investigating the root cause of the frequent non-compliance is the best course of action. It could be due to outdated software on the devices, overly strict security policies, or other factors. Once the root cause is identified, it can be addressed, which should reduce the number of quarantined devices and improve the overall security of the network.

Ignoring the issue is not a good practice. Frequent quarantining of devices indicates a potential problem with the devices' security compliance or with the NAC system's policies. It's important to investigate and address the issue to ensure the security and functionality of the network.

Disabling the automatic remediation feature may reduce the number of quarantined devices, but it also reduces the effectiveness of the NAC system. The automatic remediation feature is designed to help bring non-compliant devices into compliance, not to be disabled when it's inconvenient.

Reducing the security standards set by the NAC system may decrease the number of quarantined devices, but it also weakens the security of the network. The standards are there to protect the network and should not be reduced without a good reason.

q_nac_nonpersistent_agent_secp8

In a Network Access Control (NAC) system, a nonpersistent (or dissolvable) agent is used during the posture assessment process.

Which of the following statements about a nonpersistent NAC agent is true?

Answers:

- A nonpersistent agent remains on the device indefinitely after the posture assessment.
- A nonpersistent agent requires prior configuration on the device before it can be used.
- A nonpersistent agent cannot provide detailed information about a device's status and compliance level.
- ***A nonpersistent agent is loaded into memory during posture assessment but is not installed on the device.**

Explanation:

A nonpersistent agent is loaded into memory during the posture assessment process but is not permanently installed on the device. It is removed from memory after the assessment is completed.

A nonpersistent agent does not remain on the device after the posture assessment. It is removed from memory after the assessment is completed, which is why it is also referred to as a "dissolvable" agent.

A nonpersistent agent can be used with any device that connects to the network, including guest or IoT devices. It does not require any prior configuration on the device, making it a flexible option for assessing a variety of devices.

Even though a nonpersistent agent is not permanently installed on the device, it can still provide detailed information about the device's status and compliance level during the posture assessment process.

q_nac_plan_secp8

You are part of a committee that is meeting to define how Network Access Control (NAC) should be implemented in the organization.

Which step in the NAC process is this?

Answers:

- ***Plan**
- Apply
- Review
- Define

Explanation:

Planning is the first step in the NAC implementation process. In this step, a committee should convene and make decisions that define how NAC should work.

The third step in implementing NAC is to apply the policies. This occurs after the policies have been defined.

Review is the final step in the NAC implementation process. As business needs change, the process must be reviewed to determine whether changes are required.

Define is the second step in the NAC implementation process. After the committee has decided how NAC should work, the roles, identities, and permissions (policies) must be defined.

q_nac_vlan_01_secp8

A large enterprise recently introduced a bring your own device (BYOD) policy and is seeing an uptick in the use of Internet of Things (IoT) devices in the office.

Concerns about unauthorized network access and compliance with security standards accompany these changes.

Assess the following options and determine the MOST suitable strategy to alleviate these security concerns.

Answers:

- ***Deploy agent-based Network Access Control (NAC) with dynamic Virtual Local Area Networks (VLANs) and firewall integration.**
- Depend solely on the existing firewall for device authentication.
- Implement agentless Network Access Control (NAC) without firewall integration.
- Use Network Access Control (NAC) without employing dynamic Virtual Local Area Networks (VLANs).

Explanation:

An agent-based NAC approach combined with dynamic VLAN assignment and firewall integration provides a comprehensive solution. The solution offers robust device authentication, security policy compliance, and flexibility to accommodate various device types and user roles.

While essential for network security, firewalls do not offer the granular control and compliance enforcement that an agent-based NAC can provide.

While agentless NAC can accommodate a variety of devices without prior configuration, its information about a device's status and compliance is less detailed than that of agent-based NAC.

Employing NAC without using dynamic VLANs limits the system's ability to effectively manage network access based on user identity attributes, device type, or compliance status.

q_nac_vlan_02_secp8

An international business is experiencing an increase in remote work scenarios, resulting in a significant rise in employees using personal devices and smart appliances for work.

This development raises potential issues related to unauthorized network access and adherence to security standards.

Which of the following solutions MOST effectively addresses these security issues?

Answers:

- ***Deploy agent-based Network Access Control (NAC) with dynamic Virtual Local Area Networks (VLANs) and firewall integration.**
- Depend solely on the existing firewall for device authentication.
- Implement agentless Network Access Control without firewall integration.
- Use Network Access Control (NAC) without employing dynamic Virtual Local Area Networks (VLANs).

Explanation:

An agent-based NAC approach combined with dynamic VLAN assignment and firewall integration provides a comprehensive solution. The solution offers robust device authentication, security policy compliance, and flexibility to accommodate various device types and user roles.

While essential for network security, firewalls do not offer the granular control and compliance enforcement that an agent-based NAC can provide.

While agentless NAC can accommodate a variety of devices without prior configuration, its information about a device's status and compliance is less detailed than that of agent-based NAC.

Employing NAC without using dynamic VLANs limits the system's ability to effectively manage network access based on user identity attributes, device type, or compliance status.

q_nac_zero_trust_secp8

Which of the following BEST describes zero-trust security?

Answers:

- All devices are trusted.
- Only devices that pass authentication are trusted.
- Only devices that pass authorization are trusted.
- ***Only devices that pass both authentication and authorization are trusted.**

Explanation:

Network Access Control (NAC) is usually accomplished using a two-stage process of authentication and authorization. If the requirements for either of these stages is not met, the access request is denied. This is often referred to as zero-trust security, meaning nothing is trusted unless it can pass both the authentication and authorization stages.

Zero-trust is designed to verify trust and access of devices, filtering out those that do not authenticate properly. That being the case, not all devices are trusted when you use zero-trust security.

Zero-trust requires that devices pass both authentication and authorization--not one or the other.

5.7 Network Device Vulnerabilities

As you study this section, answer the following questions:

- For security, what is the first thing you should do when new hardware and software is turned on for the first time?
- What are the characteristics of a complex password?
- Why is it important to apply new firmware or patches for devices?
- What are major risks of hard-coded passwords on devices throughout the enterprise?
- What are the resources you can use to keep track of existing technology vulnerabilities in an organization?

In this section, you will learn to:

- Search for default passwords.
- Establish an unauthorized SSH connection.
- Secure a switch.

The key terms for this section include:

Term	Definition
Privilege escalation	A software bug or design flaw in an application that allows an attacker to gain access to system resources or additional privileges that aren't typically available.
Backdoor	An unprotected and usually lesser known access method or pathway that may allow attackers access to system resources.
Zero-day vulnerability	A software vulnerability that is unknown to the vendor that can be exploited by attackers.
Common Vulnerabilities and Exposures (CVEs)	A repository of vulnerabilities hosted by MITRE Corporation.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.3 Explain various types of vulnerabilities. • Misconfiguration • Zero-day

	2.4 Given a scenario, analyze indicators of malicious activity. • Privilege escalation 4.6 Given a scenario, implement and maintain identity and access management. • Password concepts o Password best practices ▪ Length ▪ Complexity
TestOut Security Pro	2.2 Harden Network Devices 2.2.1 Configure and Access a Switch

5.7.1 Device Vulnerabilities (Lesson Video)

Transcript:

Attackers try to exploit a wide range of potential vulnerabilities that organizations may not realize they even have. For this reason, security professionals must do all they can to reduce the attack surface their company presents. In this lesson, I'll discuss several methods for reducing your attack surface and creating a company profile that's as secure as possible. Let's first look as some often-misconfigured devices that may open up opportunities for an attack.

Remember that your attack surface increases each time a device is added to your network. This provides threat actors more opportunities to find openings in your network shield. These malicious users continually look for entry points, so it's important to ensure that each new device is properly configured and protected.
For example, firewalls are one of those entry points that are often misconfigured. Knowing this, attackers often look for firewalls with too many open ports or other security weaknesses. In addition, servers or appliances that are set up to block malicious payloads in emails could let spam and malware into the company if not configured properly. This could lead to data loss or corruption. Wireless access points are another area that, if not configured appropriately, could lead to potential attacks by allowing attackers into the production network. Using weak or outdated security protocols such as WEP is an open invitation for attackers. Once they have access to the internal network, the damage quotient is critically high.
Another attack surface to be aware of is security. The old adage says, "A chain is only as strong as its weakest link." Attackers are experts at finding the smallest hole or chink in the armor. If they want in, they'll get in. The question is how difficult it'll be for them.
Unfortunately, when something works, it's often difficult to make changes. For compatibility reasons, administrators might use outdated and insecure security protocols such as SHA1 and WEP. So every effort to utilize updated and secure protocols such as 3DES, SHA-256, and AES should be made. Another easy security attack comes from using the default credentials for your network devices, such as routers. Most attackers find that a 10 second Google query can provide the default username and password for commonplace devices. With little effort, an attacker can usually discover the make and model of the device he or she wants to attack. Those are the credentials they'll try first. To help overcome this weakness, some manufacturers are becoming more security conscious and require a change to the defaults as part of the installation process. Others make their device inoperable until the default configuration credentials are changed.
As computer networking professionals, we're very familiar with operating system and application software updates. These updates provide the latest patches, add functionality, and close security holes. The same is true for firmware. Firmware updates might include protocol updates that enable newer, more secure encryption. As an example, the famous WannaCry ransomware attack of 2017 took
advantage of a weakness inherent in the SMBv1 Protocol that allowed a remote attacker to take complete control of the affected system. This could've been prevented with security-minded updates.

Now, let's look at credential management. Unfortunately, some administrators use the same credentials in an attempt to ease the device management burden. Even though the new credentials might be different from the default ones, doing this is still a massive security risk. It might be convenient, but if any one device becomes comprised, the attacker then has access to all device configurations. Your network is greatly compromised when this happens, and you might even have to reset all your devices to factory defaults and reconfigure them.

Another potential vulnerability concern is the use of discretionary credentials. The principle of least privilege is paramount here. Users should only be given credentials with permissions that match their specific job responsibilities. If a user is given more privileges than he or she needs, an attacker could use their account to find much more potentially sensitive data than would otherwise be the case.

Next, let's go over physical access. If someone is given physical access to a device, there's very little opportunity for protection. For example, an attacker can change passwords or add credentials to a firewall with physical access. Or they can use a USB flash drive to copy data or inject malware. As such, you should confine infrastructure equipment to secure areas that only authorized personnel have access to. In addition, you should make sure that any attempted access without proper credentials sets off alarms and alerts.

One of the hard things about keeping your devices safe is trying to protect them from unknown vulnerabilities. Unknown vulnerabilities are often referred to as zero-day vulnerabilities. A zero-day vulnerability is a software security weakness that's unknown to the developer. Since these vulnerabilities are unknown, attackers who discover them quietly try to continue exploiting them without detection. Once brought to the developers' attention, they'll obviously want to fix them as quickly as possible. The best way to protect yourself from zero-day vulnerabilities is to keep your software updated. When a zero-day vulnerability is announced, patch it as quickly as possible.

Keeping up with the vastness of known vulnerabilities is also a monumental task. Fortunately, there's some help available for network administrators.

The National Institute of Standards and Technology, or NIST, has a publication that serves as a national database of known vulnerabilities. In addition, they provide vulnerability metrics called the Common Vulnerability Scoring System, or CVSS. The NIST website states that CVSS is an open framework for communicating the characteristics and severity of software vulnerabilities. After evaluating known vulnerabilities, CVSS produces a score from zero to ten, letting you calculate the vulnerability's severity. This score helps you determine the best course of action to take.

Another method of knowing about vulnerabilities is to use the Common Vulnerabilities and Exposures system, or CVE system. The system's goal is to make it easier to share information about known vulnerabilities across organizations. This is done by creating unique, standardized CVE identifiers to let you access information about specific cyber threats across multiple information sources. You can download CVE lists in several different formats. As of the writing of this lesson, there are well over 175,000 identified vulnerabilities.

That's it for this lesson. In this lesson, we discussed the need for enterprises to protect their infrastructure. Some best practices to ensure protection are to not use outdated firmware, default configurations, or weak security settings. Additionally, systems administrators need to make sure that software patches are applied in a timely manner and that bad credential management practices, like discretionary credentials, are avoided. Since securing physical access points is critical to an organization's survival, you must have an airtight infrastructure in place to do this. Finally, we talked about a few resources available to help you catalog and identify common vulnerabilities and exposures.

5.7.2 Device Vulnerability Facts

This lesson covers network device vulnerabilities.

Network Device Vulnerabilities

A knowledgeable attacker can exploit network device vulnerabilities to gain access to network resources. Network device vulnerabilities include:

Vulnerability	Description

Default accounts and passwords	Default accounts and passwords are factory defaults that already exist when a new network device is configured at installation. Default account names and passwords should be changed immediately when hardware or software is turned on for the first time.
Weak passwords	Weak passwords are passwords that are blank, too short, dictionary words, or simple. In other words, they are passwords that can be quickly identified using password-cracking tools. Password cracking is the process of recovering secret passwords from data stored in or transmitted by a computer system. Enforce complex passwords to reduce the risks of weak passwords. Complex passwords require passwords of a certain length (typically over eight characters) and a mix of character types (numbers and symbols), along with requirements that the passwords are not words, variations of words, or derivatives of the username.
Privilege escalation	Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that are not typically available to that user. Examples of privilege escalation include: • A user accessing a system with a regular user account that is able to access functions reserved for higher-level user accounts (such as administrative features). • A user who is able to access content that should only be accessible to a different user. • A user who should only have administrative access that can access content that should only be available to a regular user. Privilege escalation does not occur when a user is able to steal or hack administrator credentials and is, therefore, able to access administrative functions. Privilege escalation refers to accessing features with an account that normally should not have access to those features.
Backdoors	A backdoor is an unprotected access method or pathway. Backdoors: • Include hard-coded passwords and hidden service accounts. • Are often added during development as a shortcut to circumvent security. If they are not removed, they present a security problem. • Can be added by attackers who have gained unauthorized access to a device. When added, the backdoor can be used at a future time to bypass security controls easily. • Can be used to control the device remotely at a later date. • Rely on secrecy to maintain security. To protect against backdoors, do not allow programmers to bypass security during development. Carefully examine the code before release to remove any traces of backdoors that might have been included.
Zero-day	Zero-day vulnerabilities refer to previously unknown software or hardware flaws that attackers can exploit before developers or vendors become aware of or have a chance to fix them. The term zero-day signifies that developers have "zero days" to fix the problem once the vulnerability becomes known. These vulnerabilities are significant because they can cause widespread damage before a patch is available.

	An attacker exploiting a zero-day vulnerability can compromise systems, steal sensitive data, launch further attacks, or cause other forms of harm, often undetected. The stealth and unpredictability of zero-day attacks make them particularly dangerous. They are a favored tool of advanced threat actors, such as organized crime groups and nation-state attackers, who often use them in targeted attacks against high-value targets, such as governmental institutions and major corporations. Since these vulnerabilities are unknown to the public or the vendor during exploitation, traditional security measures like antivirus software and firewalls, which rely on known signatures or attack patterns, are often ineffective against them. The discovery of a zero-day vulnerability typically triggers a race between threat actors, who aim to exploit it, and developers, who work to patch it. Upon discovering a zero-day vulnerability, ethical security researchers usually follow a process known as responsible disclosure, designed to privately inform the vendor so a patch can be developed before the vulnerability is publicly disclosed. This practice aims to limit the potential harm caused by discovering a zero-day vulnerability. The term zero-day is usually applied to the vulnerability itself but can also refer to an attack or malware that exploits it. Zero-day vulnerabilities have significant financial value. A zero-day exploit for a mobile OS can be worth millions of dollars. Consequently, an adversary will only use a zero-day vulnerability for high-value attacks. State security and law enforcement agencies are known to stockpile zero-days to facilitate the investigation of crimes.

5.7.3 Searching for Default Passwords (Demo Video)

Transcript:

In this demonstration, we'll look at some default network device passwords. When you install a new network device, it's almost always going to have some sort of default administrative account with an associated password. You can find these by doing a search for default network device passwords, like I did up here.

Some states and countries require unique passwords for each device. Here's the website for the State of California Legislature. If I scroll down, you can see where I highlighted this line, "The preprogrammed password is unique to each device manufactured." That means in order to sell devices in this state, I have to make sure each one has a unique password. Since these sort of laws have been implemented, most manufactures are doing this for devices sold anywhere, which is a good practice. But not all devices come with unique passwords, especially enterprise equipment. On legacy equipment, having default passwords presents a security issue because these usernames and passwords are very well known and very easy to find. If someone learns your device's make and model and you haven't changed the default password, they could cause some serious damage.
Let's say we have a Netgear router, and we want to find out what the default password is. Well, I did a search, like we already saw, and here are some of the sites I found.
This first one, CIRT.net, is a pretty good site. It has over 2,100 default passwords on it. And if we come down here, we can find Netgear, and you can see a variety of different makes and models. It looks like for admin, password is a common default password for Netgear.
Let's check out another one, datarecovery.com. I'll press Ctrl+F to do a search and type in â€˜netgear'. I'll move down, and we can see a list of Netgear devices. Once again, you can see the usernames and passwords. This one's credentials are admin and password.
Let's look at one more site. This one is routerpasswords.com. I can come down and look for our Netgear devices again. I'll click on Find Passwords and see the variety of different Netgear network devices they've listed. On this site, it looks like most of the default usernames and passwords are admin and password, although there is some variance. We can also see how you interface with the device, whether it's over a web browser using HTTP, FTP, or some other way.
That's it for this demo. We searched the web for some default usernames and passwords for network devices. Remember, always change the default username and password on new network devices. Leaving default credentials presents a huge security risk.

5.7.4 Unauthorized SSH Connection (Demo Video)

Transcript:

Malicious users sometimes gain access to systems with SSH connections. It's not a good idea to allow root access through SSH, but many admins modify their systems to allow it anyway. Once a malicious user has access to the system, they can pretty much do what they want. In this demo, we're going to gain access to a Linux system and create an unauthorized SSH session.

The first thing we need to do is find our target. I'm on a Kali Linux system with a terminal open. I'm going to do a ping sweep using nmap on my subnet. I already know that my subnet is 10.10.10.0. So, to do this, I'll type '10.10.10.0/24', or I could also put in the IP range. In this case, that would be '10.10.10.1-254'. Press Enter.
My scan completes, and down here, I can see that it found four hosts. Up here, I can see that I have a device called Router with the IP of 10.10.10.1, so that's going to be my target. I'll clear the screen. And now, I want to scan just that IP to see which ports are open. The scan completes, and this, right here, is what I'm interested in. I see that SSH is running on port 22 on this system.
Now, let's use nmap and run a script that will attempt to brute force the username and password. nmap comes with something called the Nmap scripting engine. One of those scripts is for brute forcing SSH. So, let's type that in, 'nmap --script=ssh-brute 10.10.10.1', and press Enter. nmap will begin the process of attempting to brute force the username and password. nmap is using a built-in default username list and a password list to do this. First, it will grab a username and try a password. Then it goes to the next username and tries the next password. I'm on a virtual machine without a lot of system resources, so this is going to take some time. So, while this is running, I'll pause the recording.
All right. The scan finished up, and we can see, here, that it found the username Root and the password, which is a variation of the word "password." The brute force made 2,270 guesses, and it took just a little over 600 seconds to complete the process. As I mentioned, this is a virtual machine, so the scanning goes a lot slower. But I do happen to know that the password it found is near the top of the list of tens of thousands of passwords. So, depending on your system and your password file, your results could be substantially different.
Now that we know our username and password, let's try to make an unauthorized SSH connection to that system. I'm going to jump over to a Windows 10 system and try the connection from there.
All right. I'm on this Windows 10 machine, and up here, I have a program called PuTTY. PuTTY is a free and open-source terminal emulator, serial console, and network file transfer application. It supports several network protocols, including SCP, SSH, Telnet, rlogin, and other connections. It's often used to connect to switching and routing equipment. To make a connection, I'll type in the IP address for our victim machine, '10.10.10.1'. Under Port, 22 is already entered, so I'll just hit Enter.
Now, a terminal opens up. It's asking for a login name, so I'll put in 'root'. And now it wants my password, so I'll type in the variation of the word "password" that we saw earlier. Press Enter, and I'm logged in! So, whatever's possible to do in a terminal, I'm able to do--that gives me a lot of options.
Okay. I still want to look at one more thing, and that's to log in to that victim machine to see what's going on while this SSH session is running.
Let's go over to this Linux system. This is actually a Linux distribution that's used as a firewall and router. I'm going to go ahead and log in to the GUI. Once I've logged in, I'm not really able to tell that there's an SSH session going on at the exact same time.
That's it for this demo. In this demo, we performed an unauthorized SSH connection. First, we scanned our subnet with nmap to find our victim. Then we used an nmap script to brute force the username and password. We used that information to make the SSH connection to the victim machine using the program PuTTY. And we wrapped up by logging in to the victim machine's GUI to verify that the SSH session wan't detected.

5.7.5 Securing a Switch (Demo Video)

Transcript:

In this demonstration, we're going to secure access to our network switch. We're using a Cisco Small Business Switch, which has a web interface for access and configuration.

We access the switch by typing its IP address, 192.168.254.8, into a web browser. And here's the login page.
This is a new switch, so it still has the default user name and password configured of cisco and cisco, which we'll enter and then click Log In.
Now, the default user name and password for this device can be found very easily, so you can tell that this switch is not very secure right now.
Normally, the first thing you should do with a new network device is change the default user name and password.
On this device, we can use the Change Device Password link under Quick Access to do this.
This takes us directly to the User Accounts page under Administration. And here's the default user, cisco. Let's first create another user. Click Add. Enter ciscoAdmin for the user name.
Enter a strong, complex password. And we want Read/Write Management Access. And click Apply. And we can see the new ciscoAdmin user here.
We also want to change the password for the default user.
So we'll select the cisco user and click Edit. And again we want to use a complex password. And for this user, let's also change the user level. We'll select Read-Only CLI Access.
This way, if someone is able to get into the switch using the default account, they won't be able to change anything. And we'll click Apply.
So, the first thing we've done to secure access to our switch is add another admin account with a custom user name and a complex password. We then changed the default user's password and set the user level to read-only.
Doing this makes it harder for people to get into our switch here and change the configuration. And this should be the first step you do with any new network device.
Now, there's one more important step to take here. And that's to save the changes we've to the startup config. If we don't, as soon as the switch is rebooted, the changes will be lost; we don't want this.
So we'll click Save to go to the Copy/Save Configuration page. And here we want to save the contents of the Running configuration to the Startup configuration. And click Apply. Now when the device reboots, it will load our custom configuration.
That's it for this demonstration. In this demo, we secured access to our network switch. First we added a new management user with full access. We changed the default user's password. And then we saved our changes to the startup configuration file.

5.7.6 Secure a Switch (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You need to secure access to your switch, which is still configured with the default settings.

Access the switch management console through Chrome on **http://192.168.0.2** with the username **cisco** and password **cisco**.

In this lab, your task is to:

- Create a new user account with the following settings:
  - Username: **ITSwitchAdmin**
  - Password: **Admin$only1844**
  - User Level: **Read/Write Management Access (15)**
- Edit the default user account as follows:
  - Username: **cisco**
  - Password: **CLI$only1958**
  - User Level: **Read-Only CLI Access (1)**
- Save the changes to the switch's startup configuration file.

Explanation

Complete this lab as follows:

1. Log in to the CISCO switch.
    a. From the taskbar, select **Google Chrome** .
    b. In the URL field, enter **192.168.0.2** and press **Enter** .
    c. Maximize the window for better viewing.
    d. In the Username and Password fields, enter **cisco** (case sensitive).
    e. Select **Log In** .
2. Create a new user account.
    a. From Getting Started under Quick Access, select **Change Device Password** .
    b. Select **Add** .
    c. For the username, enter **ITSwitchAdmin** (case sensitive).
    d. For the password, enter **Admin$only1844** (case sensitive).
    e. For Confirm Password, enter **Admin$only1844** .
    f. For the User Level, make sure **Read/Write Management Access (15)** is selected.
    g. Select **Apply** .
    h. Select **Close** .
3. Edit the default user account.
    a. Under User Account Table, select **cisco** (the default user) and then select **Edit** .
    b. For the password, enter **CLI$only1958** .
    c. For Confirm Password, enter **CLI$only1958** .
    d. For User Level, select **Read-Only CLI Access (1)** .
    e. Select **Apply** .
4. Save the changes to the switch's startup configuration file.
    a. From the top of the switch window, select **Save** .
    b. On the left, select **Copy/Save Configuration** .
    c. On the right, under *Source File Name* , make sure **Running configuration** is selected.
    d. Under *Destination File Name* , make sure **Startup configuration** is selected.
    e. Select **Apply** .
    f. Select **OK** .
    g. Select **Done** .

5.7.7 Practice Questions (Section Quiz)

q_dev_vuln_backdoor_01_secp8

While developing a network application, a programmer adds functionally that allows her to access the running program without authentication so she can capture debugging data. The programmer forgets to remove this functionality prior to finalizing the code and shipping the application.

Which type of security weakness does this describe?

Answers:

- ***Backdoor**
- Privilege escalation
- Weak password
- Buffer overflow

Explanation:

A backdoor is an unprotected access method or pathway. Backdoors may include hard-coded passwords or hidden service accounts. They are often added during development as a shortcut to circumvent security. If they are not removed, they present a security problem.

Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that would typically not be available to the user.

Weak passwords are passwords that are blank, too short, dictionary words, or not complex enough. This allows them to be quickly identified using password-cracking tools.

A buffer overflow occurs when the operating system or an application does not properly enforce boundaries for how much and which type of data can be inputted.

q_dev_vuln_backdoor_02_secp8

An attacker was able to gain unauthorized access to a mobile phone and install a Trojan horse so that he or she could bypass security controls and reconnect later.

Which type of attack is this an example of?

Answers:

- Social engineering
- Replay
- ***Backdoor**
- Privilege escalation

Explanation:

A backdoor is an unprotected access method or pathway. Backdoors:

- Include hard-coded passwords and hidden service accounts.
- Are often added during development as a shortcut to circumvent security. If they are not removed, they present a security problem.
- Can be added by attackers who have gained unauthorized access to a device. When added, the backdoor can be used at a future time to easily bypass security controls.
- Can be used to remotely control the device at a later date.
- Rely on secrecy to maintain security.

Social engineering attacks involve stealing information or convincing someone to perform an inappropriate activity via email, via phone, or in person.

A replay attack is a network attack that occurs when an attacker intercepts data and fraudulently delays or re-transmits it.

Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that aren't typically available to that user.

q_dev_vuln_backdoor_03_secp8

In an effort to increase the security of your organization, programmers have been informed they can no longer bypass security during development.

Which vulnerability are you attempting to prevent?

Answers:

- Social engineering
- Replay
- ***Backdoor**
- Privilege escalation

Explanation:

A backdoor is an unprotected access method or pathway. Backdoors:

- Include hard-coded passwords and hidden service accounts.
- Are often added during development as a shortcut to circumvent security. If they are not removed, they present a security problem.
- Can be added by attackers who have gained unauthorized access to a device. When added, the backdoor can be used at a future time to easily bypass security controls.
- Can be used to remotely control the device at a later date.
- Rely on secrecy to maintain security.

Social engineering attacks involve stealing information or convincing someone to perform an inappropriate activity via email, phone, or in person.

A replay attack is a network attack that occurs when an attacker intercepts data and fraudulently delays or re-transmits it.

Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that aren't typically available to that user.

q_dev_vuln_complex_secp8

Which of the following are characteristics of a complex password? (Select two.)

Answers:

- Has a minimum of six characters
- ***Has a minimum of eight characters**
- Consists of letters and numbers only
- ***Consists of letters, numbers, and symbols**
- Has a maximum of fifteen characters

Explanation:

Complex passwords require a certain length (typically over eight characters) and a mix of character types (numbers and symbols) along with requirements that the password not consist of words, variations of words, or derivatives of the username.

There is no maximum character limit for a complex password.

q_dev_vuln_cracking_secp8

An attacker has gained access to the administrator's login credentials.

Which type of attack has MOST likely occurred?

Answers:

- Privilege escalation
- Backdoor
- ***Password cracking**
- Buffer overflow

Explanation:

Password cracking is the process of recovering secret passwords from data that has been stored in or transmitted by a computer system. If an attacker has gained access to the administrator's login credentials, this is most likely the cause of a password-cracking attack.

Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that would typically not be available to the user.

A backdoor is an unprotected access method or pathway. Backdoors may include hard-coded passwords or hidden service accounts.

A buffer overflow attack occurs when the operating system or an application does not properly enforce boundaries for how much and which type of data can be inputted.

q_dev_vuln_privilege_01_secp8

A relatively new employee in the data entry cubical farm was assigned a user account similar to the other data entry employees' accounts. However, audit logs have shown that this user account has been used to change ACLs on several confidential files and has accessed data in restricted areas.

This situation indicates which of the following has occurred?

Answers:

- Physical security
- Social engineering
- External attack
- ***Privilege escalation**

Explanation:

This situation describes the result of a successful privilege escalation attack. If a low-end user account is detected performing high-level activities, it is obvious that the user account has somehow gained additional privileges.

Physical security is the protection of corporate assets from threats such as theft or damage.

Social engineering attacks involve stealing information or convincing someone to perform an inappropriate activity via email, phone, or in person.

External attacks are when unauthorized individuals try to breach a network from off-site.

q_dev_vuln_privilege_02_secp8

An attacker has obtained the logon credentials for a regular user on your network.

Which type of security threat exists if this user account is used to perform administrative functions?

Answers:

- ***Privilege escalation**
- Social engineering
- Replay
- Impersonation

Explanation:

Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that are typically not available to normal users. Examples of privilege escalation include:

- A user accessing a system with a regular user account and successfully accessing functions reserved for higher-level user accounts (such as administrative features).
- A user who is able to access content that should be accessible only to a different user.
- A user who should have only administrative access being able to access content that should only be accessible to a regular user.

Privilege escalation does not occur when a user is able to steal or hack administrator credentials and is, therefore, able to access administrative functions. Privilege escalation refers to accessing features with an account that normally should not have access to those features.

q_dev_vuln_privilege_03_secp8

Travis and Craig are both standard users on the network. Each user has a folder on the network server that only they can access. Recently, Travis has been able to access Craig's folder.

This situation indicates which of the following has occurred?

Answers:

- Social engineering
- Replay
- External attack
- ***Privilege escalation**

Explanation:

This situation describes the result of a successful privilege escalation attack. If a user is able to access content that should only be accessible to a different user, it is obvious that a privilege escalation attack has occurred.

Social engineering attacks involve stealing information or convincing someone to perform an inappropriate activity via email, phone, or in person.

A replay attack is a network attack that occurs when an attacker intercepts data and fraudulently delays or re-transmits it.

External attacks are when unauthorized individuals try to breach a network from off-site.

q_dev_vuln_zero-day_01_secp8

In a rapidly evolving IT environment, a cloud service provider offers various services to businesses, enabling them to store and process data securely. To enhance security, the provider regularly updates its systems and software.

Despite these efforts, a security researcher discovers a previously unknown vulnerability in one of the cloud-specific applications, leaving customer data exposed to potential threats.

In this scenario, which vulnerability is the security researcher likely to have found in the cloud-specific application?

Answers:

- ***Zero-day vulnerability**
- Network misconfiguration
- SQL injection vulnerability
- Cross-site scripting (XSS) vulnerability

Explanation:

A zero-day vulnerability refers to a security flaw previously unknown to the vendor or developer. The discovery of this previously unknown vulnerability aligns with the characteristics of a zero-day vulnerability.

In the scenario provided, the focus is on a vulnerability identified within a cloud-specific application that is distinct from a misconfiguration in the network settings and often involves errors in firewall rules, access controls, or other network-related configurations.

SQL injection is a prevalent type of cyberattack where malicious actors exploit vulnerabilities in web applications to insert malicious SQL queries into the application's database.

Malicious script injection, also known as cross-site scripting (XSS), is when attackers inject malicious scripts into web pages to compromise users' sessions or steal sensitive information.

q_dev_vuln_zero-day_02_secp8

In the context of information security, an organization discovers a zero-day vulnerability in its database software.

At the same time, a known hacking group has expressed intentions to target entities using this specific software.

Which of the following BEST describes this situation's relation to vulnerability, threat, and risk?

Answers:

- The organization conducts regular vulnerability assessments to maintain its security posture.
- The organization mitigates the risk by improving physical security and firewall configurations.
- The organization hires an external cybersecurity firm to identify potential threats.
- ***The organization increases its risk of a security breach due to the threat and vulnerability.**

Explanation:

This option illustrates a scenario in which an external group (threat) threatens a vulnerability (the software weakness), raising the possibility of a security breach (risk).

While conducting regular vulnerability assessments is a good practice, it does not demonstrate the relationship between vulnerability, threat, and risk.

Focusing on physical security and firewall configurations focuses on risk mitigation techniques but does not illustrate the interplay between vulnerability, threat, and risk.

While hiring external cybersecurity expertise can be part of a comprehensive security strategy, this choice does not demonstrate the relationship between vulnerability, threat, and risk.

q_dev_vuln_zero-day_03_secp8

A major software vendor becomes aware of a new zero-day vulnerability in one of its products due to an anonymous tip. The vulnerability could potentially allow unauthorized access to sensitive data stored in the software.

The vendor is currently creating a patch to address the issue.

Which of the following BEST describes the current risk to the software users and the appropriate response from the software vendor?

Answers:

- Since the vendor knows about the vulnerability, there is minimal risk. The vendor should alert all users about the vulnerability immediately and provide mitigation steps.
- ***The risk to the users is significant, and the vendor should quietly create a patch without informing the users until it is ready.**
- There is no risk to users as long as the vendor does not disclose the vulnerability. The vendor should continue its usual operations without interruption.
- The risk to users is unknown, and the vendor should contact individual users to assess potential damage before proceeding.

Explanation:

Zero-day vulnerabilities represent significant risk, and the vendor should prioritize creating a patch. Disclosing the vulnerability to the public before a patch is ready could increase the risk.

Alerting users about the vulnerability could also alert potential threat actors, increasing the risk before a patch is ready.

Even if the vendor does not disclose the vulnerability, there is a risk while the vulnerability exists, especially since someone is aware of it, as indicated by the anonymous tip.

The risk is significant due to the nature of the vulnerability, and contacting individual users is not feasible, nor will it reduce the risk. The vendor should focus on creating a patch.

q_dev_vuln_zero-day_04_secp8

An information security analyst at a tech company reviews a security report outlining recent attack vectors against the company's systems.

The analyst identifies potential risks related to unpatched software vulnerabilities still unknown to the vendor and risks associated with weak cryptographic algorithms. The analyst wants to prioritize these risks to decide on immediate remedial action.

Based on the provided scenario, what BEST describes an unknown vulnerability in software that the vendor has yet to discover or patch, and that attackers are actively exploiting?

Answers:

- ***Zero-day vulnerability**
- On-path attack
- Rainbow table attack
- Public key infrastructure flaw

Explanation:

The term "zero-day" describes a software vulnerability that the vendor has yet to identify. Attackers exploit this vulnerability and jeopardize systems.

In an on-path attack, an attacker intercepts and might alter communication between two parties. However, this threat does not connect to a software vulnerability unpatched and unknown by the vendor.

Attackers crack encrypted passwords using this method, referencing a precomputed table of hash values for possible passwords. This method does not address a software vulnerability that the vendor has not identified or patched.

While issues in public key infrastructure (PKI) introduce vulnerabilities into cryptographic systems, they do not point to an unpatched software flaw unknown to the vendor.

5.8 Network Applications

As you study this section, answer the following questions:

- How does application vulnerability scanning differ from general vulnerability scanning?
- What is package monitoring used for?
- What security measures should you incorporate to control the use of networking software?
- What is static analysis?
- What is dynamic analysis?

In this section, you will learn to:

- Configure application control software.

The key terms for this section include:

Term	Definition
Peer-to-peer (P2P) software	Software that allows users to share content without centralized servers or centralized access control.
Instant messaging	Real-time text messaging communication that supports picture, music, and document exchange.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	4.3 Explain various activities associated with vulnerability management.
	• Identification methods
	  ○ Application security
	  ▪ Static analysis
	  ▪ Dynamic analysis
	  ▪ Package monitoring

TestOut Security Pro	3.2 Implement Application Defenses 3.2.1 Implement Application Whitelisting 3.2.3 Configure web application security

5.8.1 Network Application Security (Lesson Video)

Transcript:

In this lesson, we'll talk about how applications can pose a threat to your network and how to manage those risks. We start by looking for vulnerabilities.

Vulnerability scanning is a crucial aspect of cybersecurity that aids in identifying weaknesses in a system's security defenses.
Imagine your organization's network and software systems as a secure facility and vulnerability scanning as a diligent guard
searching for weak points in its fence.
To achieve this, cybersecurity professionals use specialized tools like openVAS and Nessus. These tools act as vigilant security guards, examining various aspects of your network and systems. They check for issues such as outdated software, misconfigured settings, and missing patches, which can serve as potential entry points for malicious actors.
However, when it comes to application security, a more tailored approach is needed. Applications are like unique rooms within your secure facility, each with its own vulnerabilities. Traditional vulnerability scanning tools can't fully grasp these intricacies. For this, we turn to specialized application scanners, penetration testing frameworks, and code testing methods.
Application vulnerability scanning is a refined process that zeros in on the specific weaknesses within software applications. Think of it as having an expert locksmith who knows the ins and outs of each room's lock and key.
This scanning method employs two techniques: static analysis and dynamic analysis. Static analysis involves reviewing the application's code without executing it, much like scrutinizing the blueprints of a room. Dynamic analysis, on the other hand, tests the running applications, uncovering issues like unvalidated inputs, broken access controls, and SQL injection vulnerabilities by actively trying to breach the application's defenses.
These vulnerabilities are unique to software applications, and their discovery requires specialized tools and techniques.
Managing these risks necessitates a separate approach from general vulnerability scanning
because applications often follow their own update cycles and have distinct security challenges.
In the complex world of software development, applications often rely on external components, like libraries and dependencies. Imagine these components as pieces of machinery inside your secure facility.
Package monitoring is the process of keeping a watchful eye on these components. It ensures they're up-to-date and free from known vulnerabilities that hackers might exploit. This is essential because just
like a compromised piece of machinery can weaken the security of your facility, outdated or vulnerable software components can jeopardize your application's security.
To achieve this, organizations use automated tools like Software Composition Analysis tools. These tools constantly compare your software inventory against databases of known vulnerabilities, ensuring that your components are sturdy and reliable. They can also suggest updates or replacements when vulnerabilities are discovered.

In addition to automated tools, organizations establish governance policies around software usage. These policies mandate regular audits of software packages, require approvals for new additions, and outline procedures for updating or patching software in case vulnerabilities are identified.
So, think of vulnerability scanning, application vulnerability scanning, and package monitoring as the various layers of security in your facility: one guarding the perimeter, another securing the rooms, and the last ensuring the reliability of the machinery within. Together, they form a comprehensive defense against modern cyber threats.

5.8.2 Configure Application Control Software (Demo Video)

Transcript:

In this demonstration, we're going to look at application control. Application controls focus on controlling the applications a user at a workstation can access over a network. This is similar to what a firewall does.

For example, a packet filtering firewall examines traffic going through it. Based on its rule set, a firewall will block traffic using specific ports or protocols, so you should be able to control application access with a firewall.
However, if you have a savvy user trying to access a forbidden application on the internet, the user can reconfigure the session to use a protocol or port that is commonly left open on most network firewalls, for example, port 80.
A user could reconfigure a game application blocked by the firewall to use port 80. The packet filtering firewall doesn't examine the contents of the packet, only the protocol--TCP, UDP--and the port.
So, a firewall wouldn't block the game traffic redirected to a port that's used for legitimate traffic.
Application control software, on the other hand, focuses on the application signature.
In much the same way that antivirus software uses signatures to detect viruses on your system, application control software uses application signatures to identify traffic associated with a specific application.
It disregards the port number used.
In this demonstration, we're going to configure application control software to block specific network applications.
To do this, we have our network set up similar to what is shown in this picture here, with user computers accessing the network through a shared switch--some wirelessly, some wired--and ultimately passing through the firewall that also has application control software configured on it, to our router, and then out to the internet.
In other words, all traffic passes through here; and as such, the application control software can look at those packets before it allows entry or exit from the network.
For this demo, we'll be using the untangle software that includes an application control module. Certain aspects of this program are free. You have to pay for the application control, but you can try it out for free for 14 days.
You can click up here to see what applications are currently installed. You can see there's a variety of applications or modules that they have.
We have captive portal, which requires, for instance, users to accept a use agreement before getting onto our network; there's a firewall; there are some phishing blockers or virus protection.
We're interested in this application control. Even though the application control is up and running and turned on, it's actually not filtering much. Before we can use it, we need to tell it which applications we would like to filter.
In order to demonstrate this, before I configure that, I want to show you that I have access on a machine that's within my network. I'm going to switch over to my Windows 10 machine now.
My Windows 10 machine is one of these computers over here. Let's go look at it now. As you can see, my Windows 10 machine has access to the internet right now.
We can go out and do a Google search on application controllers--or, in this case, a Bing search--and you can see it returns results.
Perhaps I want to check Facebook and see what's going on in my life, or I could pull up Skype and start to chat my life away on Skype. You can see that these applications are working.
Let's go ahead now and switch back to our application control and enable some rules to prohibit access to these services. We'll come here to the Applications tab, and then we need to go down and find those specific services.
You can see there's a wide variety of services that are listed here. I wish there was a better way to search for the particular service that you're wanting to block. For now, we'll have to be content with just scrolling down. Look at the thousands, over 2000 different applications, that you can block.
So, I find Facebook here and I'm going to go ahead and block that. I could also flag it so that it ends up in one of my reports.

This Tarpit option right here, simply allows the connection to be established, but then the application control 'software untangle' drops the information so that both sides think there's a connection, but no data actually gets through. We blocked Facebook.
Let's go ahead and scroll down and block Skype now. There's lots of different Skype services. I'm going to just block them all. You can see that it's fairly granular in what you can do, allow or prohibit.
Once you're done, then you need to click the Save button. Unfortunately, the application doesn't scale well, so you'll have to scroll down to get to that Save button. There it is right there.
Once you've clicked that Save button, now those applications and those rules will be enforced. Let me now switch back over to my Windows 10 environment and see the impact of those rule changes.
Let's go ahead and try to refresh Facebook. You'll notice that nothing works. As we refresh Facebook, you see that the page can't be reached. However, other pages are still accessible, and we can still do our search.
Let's go investigate Skype now. You can see here that my contacts are no longer online, and it's having trouble connecting over there.
Again, you can get pretty granular on what you want to block and what you want to allow. Let's switch back over to the untangle server now. You can see how effective that was in blocking the traffic.
We can go look and it monitors how many sessions are there. Looks like it had 2500 scans and allowed 1700 of those. It flagged 70 different applications, kept track of those.
You have kind of a dashboard around what it has done. You can easily turn this on and off by simply clicking right here.
And then if we switch back over to our Windows 10, you'll see that all of a sudden I can now access Skype and Facebook. Let's do that. No problem. And if we pull up Skype, you can see that it's back online.
Switching back over to our untangle server', let's just look at a couple more things.
This first column is just the application signature. You can see the name, the category.
This is nice because you can ban whole groups of applications based on the type of category, or the presumed productivity associated with that application, risk associated with it, and then a short description of that. There are others.
You can go look at the specific roles, or you can view reports behind this. So, once again, very easy to set up, but very useful in helping to control your environment, and what's going on within your enterprise.
That's it for this demonstration. In this demo, we introduced you to application control software. We first talked about what application control software does.
We then configured the untangle application control software to block specific traffic--in this case, Skype and Facebook.
Then we went to the client system on the internal network and we tried to access a forbidden application and it was blocked.

5.8.3 Network Application Facts

This lesson covers the following topics:

- Vulnerability scanning
- Application vulnerability scanning
- Package monitoring
- Managing networking software

Vulnerability Scanning

Vulnerability scanning supports application security, as it helps to locate and identify misconfigurations and missing patches in software. Advanced vulnerability scanning techniques focused on application security include specialized application scanners, pen-testing frameworks, and static and dynamic code testing.

Vulnerability scanning tools like openVAS and Nessus are popular tools offering a broad range of features designed to analyze network equipment, operating systems, databases, patch compliance, configuration, and many other systems. While these tools are very effective, application security analysis warrants much more specialized approaches. Several specialized tools exist to more deeply analyze how applications are designed to operate and can locate vulnerabilities not typically identified using generalized scanning approaches.

Application Vulnerability Scanning

Application vulnerability scanning describes a specialized vulnerability scanning method for identifying software application weaknesses. This includes static analysis (reviewing application code without executing it) and dynamic analysis (testing running applications), which can identify issues like unvalidated inputs, broken access controls, and SQL injection vulnerabilities. Application vulnerability scanning is typically handled separately from general vulnerability scanning due to the unique nature of software applications and the specific types of vulnerabilities they introduce. General vulnerability scanning is designed to detect system-wide or network-wide weaknesses, such as out-of-date software or misconfigured firewalls.

In contrast, application vulnerability scanning evaluates the coding and behavior of individual software applications. It looks for issues like cross-site scripting (XSS), SQL injection, and insecure direct object references unique to software applications. These application-specific vulnerabilities require specialized tools and techniques to identify and mitigate and are generally different from the scanning tools used in general vulnerability scanning. Applications frequently have their own release and update cycles, separate from the rest of the environment, necessitating a more targeted vulnerability management process.

Package Monitoring

Another important capability in application vulnerability assessment practices includes package monitoring. Package monitoring is associated with vulnerability identification because it tracks and assesses the security of third-party software packages, libraries, and dependencies used within an organization to ensure that they are up-to-date and free from known vulnerabilities that malicious actors could exploit. Package monitoring is associated with the management of software bill of materials (SBOM) and software supply chain risk management practices.

In an enterprise setting, package monitoring is typically achieved through automated tools and governance policies. Automated software composition analysis (SCA) tools track and monitor the software packages, libraries, and dependencies used in an organization's codebase. These tools can automatically identify outdated packages or packages with known vulnerabilities and suggest updates or replacements. They work by continuously comparing the organization's software inventory against various databases of known vulnerabilities, such as the National Vulnerability Database (NVD) or vendor-specific advisories.

In addition to these tools, organizations often implement governance policies around software usage. These policies may require regular audits of software packages, approval processes for adding new packages or libraries, and procedures for updating or patching software when vulnerabilities are identified.

Managing Networking Software

Use the following to control the use of networking software:

- Have a written policy that identifies the allowed or prohibited usage of all software.
- Use Group Policy or other methods to prevent installation of the software.
- Block firewall ports that are used by the software.
- Consider implementing an application control solution.
    - A firewall alone may be insufficient in blocking the use of network applications.
        - Knowledgeable users can circumvent firewall ACLs by reconfiguring network applications to use ports commonly left open.
        - Packet filtering firewalls do not inspect the contents of a packet. Only the source IP address, destination IP address, protocol, and port are used to determine if a packet should be blocked.
    - An application control solution can be used to block unauthorized network applications.
        - Application control implementations use application signatures to identify specific applications.
        - The contents of packets are inspected and compared against these signatures to identify the associated application.
    - An application allow list is defined centrally and applied to all network devices.
        - Only applications contained in the allow list are allowed.

        - Several actions can be applied to applications that are not on the allow list:
            - Flagged applications are allowed, but a violation is logged when they are identified.
            - Blocked applications are not allowed and are blocked. The session is dropped if it uses UDP and reset if it uses TCP.
            - Tarpitted applications are not allowed. However, the connection between hosts is kept alive while the application data itself is silently dropped. This makes it appear to both hosts that the other host is receiving the data but not responding.

              Not all application control solutions support tarpitting application traffic.

    - If a user tries to use a disallowed application, they can be prompted to contact the help desk or system administrator to get the application reviewed and approved for use.

5.8.4 Practice Questions (Section Quiz)

q_net_app_allow_list_secp8

As the IT manager of a company, you've decided to implement an application allow list to enhance network security and control software usage.

What is the MOST effective way to proceed?

Answers:

- Implement the allow list immediately and block all applications not on the list without informing the employees.
- Create the allow list based solely on the IT department's preferences without consulting other departments.
- ***Implement the allow list, inform employees about the change, and provide a process for requesting additions to the list.**
- Implement the allow list but allow employees to install and use applications not on the list if they feel it's necessary.
- Create an allow list but do not enforce it, leaving it as a guideline for employees.

Explanation:

Implement the allow list, inform employees about the change, and provide a process for requesting additions to the list is the most effective as it ensures that the allow list is implemented and adhered to, while also providing a clear communication channel for employees. By allowing employees to request additions to the list, you ensure that necessary software isn't inadvertently blocked, and you maintain a level of flexibility.

Implementing the allow list immediately and blocking all applications not on the list without informing the employees could lead to confusion and frustration among employees. It's important to communicate changes and provide clear guidelines.

Creating the allow list based solely on the IT department's preferences without consulting other departments could lead to essential software being blocked, hindering productivity. It's important to consult with all departments to ensure necessary software is included in the allow list.

Implementing the allow list but allowing employees to install and use applications not on the list if they feel it's necessary undermines the purpose of the allow list and could lead to security vulnerabilities.

Creating an allow list but not enforcing it, and leaving it as a guideline for employees does not effectively control software usage and could lead to continued unauthorized software usage and potential security risks.

q_net_app_monitor_software_usage_secp8

As a company grows, so does its attack surface and the desirability for a malicious actor to compromise its systems. A company must monitor all software usage, secure applications, third-party software, libraries, and dependencies to keep systems secure.

What are some ways to BEST accomplish this? (Select two.)

Answers:

- ***Implementing application vulnerability scanning**
- ***Using package monitoring**
- Ignoring software updates and patches
- Relying solely on firewall protections
- Allowing all employees unrestricted access to install and use any software

Explanation:

The following are the best answers (most correct):

- Application vulnerability scanning is a specialized method for identifying weaknesses in software applications. It includes static analysis, which involves reviewing application code without executing it, and dynamic analysis, which involves testing running applications. These methods can identify issues such as unvalidated inputs, broken access controls, and SQL injection vulnerabilities. By identifying these vulnerabilities, a company can take steps to address them and thereby enhance its system security.
- Package monitoring involves tracking and assessing the security of third-party software packages, libraries, and dependencies. Automated software composition analysis (SCA) tools can be used for this purpose. These tools can identify outdated packages or packages with known vulnerabilities, and suggest updates or replacements. By keeping these elements up-to-date and free from known vulnerabilities, a company can reduce its attack surface and enhance its system security.

The following are not the best (or incorrect) answers:

- Ignoring software updates and patches is incorrect because software updates and patches often address known vulnerabilities. Ignoring them means leaving these vulnerabilities unaddressed, which increases the company's attack surface and makes its systems more susceptible to compromise.
- While firewalls are an important part of system security, relying solely on them is not sufficient. Firewalls primarily control incoming and outgoing network traffic based on predetermined security rules, but they do not inspect the contents of a packet. Therefore, they cannot identify specific applications or block unauthorized network applications. For this reason, application control solutions, which use application signatures to identify specific applications, are also necessary.
- Allowing all employees unrestricted access to install and use any software is incorrect because not all software is secure. Allowing employees unrestricted access to install and use any software increases the risk of installing software with known vulnerabilities or malicious software. Therefore, it is important to have approval processes and regular audits of software packages to ensure that only secure and necessary software is installed and used.

q_net_app_package_tracking_secp8

A system administrator at a software development company is working on integrating package monitoring into the organization's vulnerability management strategy. The administrator aims to track software packages and applications to ensure they remain free from vulnerabilities and continue to support the firm's security framework.

As the system administrator incorporates package monitoring into the vulnerability management process, which actions will MOST likely get prioritized to enhance the effectiveness of this approach? (Select two.)

Answers:

- ***Tracking outdated software packages**
- Manually updating software every day
- ***Monitoring software repositories for new updates**
- Buying the latest antivirus software every month
- Block firewall ports that are used by the software.

Explanation:

Keeping track of outdated software packages is vital since they can often become targets for attackers if they contain vulnerabilities that newer versions have fixed.

Regularly checking software repositories helps administrators stay informed about the latest updates and patches, ensuring the organization's software remains current and less susceptible to known vulnerabilities.

While regular software updates are crucial, manual daily updates can be excessive and may disrupt operations. Instead, administrators should set a balanced update schedule based on the software's criticality and vulnerabilities.

Regularly updating antivirus definitions is crucial, but purchasing entirely new antivirus software every month is unnecessary and not directly related to package monitoring in vulnerability management.

While blocking firewall ports is an important part of managing networking software, it has no direct impacting on package monitoring.

q_net_app_patch_secp8

A cybersecurity analyst for a large organization permits employees to use instant messaging (IM) services on their devices. Despite using encryption, the analyst's concern is the potential software vulnerabilities and difficulty scanning messages and attachments for threats. Which actions should the cybersecurity analyst use to address this concern?

Answers:

- ***Regularly update and patch the instant messaging apps to address any known software vulnerabilities.**
- Disable all instant messaging services on Windows, Android, and iOS devices to prevent any potential security risks.
- Implement additional encryption layers on top of the existing instant messaging (IM) services to enhance security further.
- Allow employees to use instant messaging services without any changes since the encryption already provides sufficient security.

Explanation:

The recommended action is regularly updating and patching the instant messaging apps to address any known software vulnerabilities. While encryption offers added security, keeping the apps up-to-date is crucial to mitigating potential risks.

Disabling all instant messaging services may hinder employee communication and collaboration, possibly affecting productivity.

Implementing additional encryption layers may introduce complexity and compatibility issues, and the existing encryption already provides a considerable amount of security.

Assuming that the encryption alone provides sufficient security may overlook the importance of addressing potential software vulnerabilities through regular updates and patches.

q_net_app_sbom_sca_secp8

Which of the following are important practices in application vulnerability assessment to ensure the security of third-party software packages, libraries, and dependencies used within an organization?

Answers:

- ***SBOM**
- ***SCA**
- VPN
- SSL
- DHCP

Explanation:

The following are the correct answers:

- SBOM stands for Software Bill of Materials. It is a list of components in a piece of software. SBOMs contribute to software transparency and allow for the identification and management of open source components and third-party software dependencies, which in turn helps to identify known vulnerabilities and manage software supply chain risks.
- SCA stands for Software Composition Analysis. It is a method used to identify open source components and third-party software dependencies in a codebase, and to detect known vulnerabilities in them. By continuously comparing the organization's software inventory against various databases of known vulnerabilities, SCA tools can suggest updates or replacements for outdated packages or packages with known vulnerabilities.

VPN stands for Virtual Private Network. While VPNs are important for securing internet connections and protecting data in transit, they do not directly contribute to the management of third-party software packages, libraries, and dependencies.

SSL stands for Secure Sockets Layer. It is a protocol for establishing encrypted links between a web server and a browser in online communications. While SSL is important for securing data in transit, it does not directly contribute to the management of third-party software packages, libraries, and dependencies.

DHCP stands for Dynamic Host Configuration Protocol. It is a network protocol that enables a server to automatically assign an IP address to a computer from a defined range of numbers configured for a given network. While DHCP is important for managing network configurations, it does not directly contribute to the management of third-party software packages, libraries, and dependencies.

q_net_app_signature_secp8

What do application control solutions use to identify specific applications?

Answers:

- Flags
- Whitelists
- Packet inspection
- ***Application signatures**

Explanation:

Application control implementations use application signatures to identify specific applications.

Group Policy is used to define security policies on a Windows operating system. Application control systems do not use group policies to identify specific applications.

Whitelists are used to define which applications are allowed on network devices.

Packet inspection is performed by firewalls, not application control solutions.

q_net_app_tarpit_secp8

You have implemented a new application control solution. After monitoring traffic and use for a while, you have noticed an application that continuously circumvents blocking.

How should you configure the application control software to handle this application?

Answers:

- Block
- Flag
- ***Tarpit**
- Drop

Explanation:

When using tarpit, the connection between hosts is kept alive while the application data itself is silently dropped. This makes it appear to both hosts that the other host is receiving the data but is not responding. Some malicious applications notice they are being blocked and circumvent the issue. Using tarpit prevents the application from realizing it has been blocked and stops it from circumventing security controls.

Blocked applications are not allowed and are blocked. The session is dropped if it uses UDP and is reset if it uses TCP.

Flagged applications are allowed, but a violation is logged when they are identified.

Drop is not a application control software configuration option.

q_net_app_training_secp8

You are the IT manager of a company that develops its own software applications. You've been tasked with enhancing the security of these applications. You decide to implement application vulnerability scanning.

What is the MOST effective way to proceed?

Answers:

- Implement application vulnerability scanning only for applications that have previously been compromised.
- Implement application vulnerability scanning without informing the development team.
- ***Implement application vulnerability scanning and provide comprehensive training to the development team on the process and how to address identified vulnerabilities.**
- Implement application vulnerability scanning but ignore minor vulnerabilities identified.
- Implement application vulnerability scanning and immediately halt all development activities until all identified vulnerabilities are addressed.

Explanation:

Implementing application vulnerability scanning and providing comprehensive training to the development team on the process and how to address identified vulnerabilities is the most effective way to proceed as it not only introduces a method to identify vulnerabilities but also ensures that the development team is equipped to address these vulnerabilities. Training and clear communication are key to ensuring that the solution is used effectively and that vulnerabilities are properly addressed.

Implementing application vulnerability scanning only for applications that have previously been compromised could leave other applications exposed to potential vulnerabilities. It's important to scan all applications to ensure comprehensive security.

Implementing application vulnerability scanning without informing the development team could lead to confusion and resistance among the development team. It's important to communicate changes and provide clear guidelines.

Implementing application vulnerability scanning but ignoring minor vulnerabilities identified could leave the applications exposed to potential security risks. Even minor vulnerabilities can be exploited and should be addressed.

Implementing application vulnerability scanning and immediately halting all development activities until all identified vulnerabilities are addressed is not very effective. While addressing vulnerabilities is important, halting all development activities could disrupt the company's operations. It's important to find a balance between addressing vulnerabilities and maintaining productivity.

q_net_app_vuln_scan_secp8

Which of the following security actions represents a non-intrusive scanning type of framework?

Answers:

- ***Vulnerability scanning**
- Penetration testing
- Access control list (ACL)
- Keylogger

Explanation:

Whether they use purely passive techniques or some sort of active session or agent, vulnerability scanners represent a non-intrusive scanning type. The scanner identifies vulnerabilities in its database by analyzing things such as build and patch levels or system policies.

Penetration testing that uses exploitation frameworks are "active" and "intrusive" and detectable.

An ACL is a list of permissions associated with a network object, such as a router or a switch, that controls traffic at a network interface level. An ACL is not a type of framework.

A keylogger is spyware that actively attempts to steal confidential information by recording keystrokes. The attacker will usually hope to discover passwords or credit card data. It is not a framework in itself.

q_net_app_white_secp8

You are implementing a new application control solution.

Prior to enforcing your application whitelist, you want to monitor user traffic for a period of time to discover user behaviors and log violations for later review.

How should you configure the application control software to handle applications not contained in the whitelist?

Answers:

- Block
- ***Flag**
- Tarpit
- Drop

Explanation:

When using an application control solution, an application whitelist is defined centrally and applied to all network devices. Only applications contained in the whitelist are allowed. Applications not whitelisted can have several actions applied:

- Blocked applications are not allowed. The session is dropped if it uses UDP and reset if it uses TCP.
- Flagged applications are allowed, but a violation is logged when they are identified.
- Tarpitted applications are not allowed. However, the connection between hosts is kept alive while the application data itself is silently dropped. This makes it appear to both hosts that the other host is receiving the data but not responding.

Not all application control solutions support tarpitting application traffic.

5.9 Switch Security and Attacks

As you study this section, answer the following questions:

- How are switches indirectly involved in Address Resolution Protocol (ARP) poisoning?
- How does the attacker hide his identity when performing media access control (MAC) address spoofing?
- What is the function of a trunk port?
- How is port security different from port filtering?
- What are some important considerations when deciding which secure protocols to implement?

In this section, you will learn to:

- Harden a switch.
- Secure access to a switch.
- Use best practices to ensure switch security.

The key terms for this section include:

Term	Definition
Virtual LAN (VLAN)	A logical grouping of computers based on switch port.
MAC filtering/port security	A switch feature that restricts connection to a given port based on the MAC address.
Port authentication	A switch feature that follows the 802.1x protocol to allow only authenticated devices to connect.
Content-addressable memory (CAM) table	A table maintained by a switch that contains MAC addresses and their corresponding port locations.
Dynamic Host Configuration protocol (DHCP) snooping	A security feature on some switches that filters out untrusted DHCP messages.
Dynamic ARP Inspection (DAI)	A security feature on some switches that verifies each ARP request has a valid IP to MAC binding.
MAC flooding	An attack that overloads a switch's MAC forwarding table to make the switch function like a hub.
ARP spoofing	An attack in which the attacker's MAC address is associated with the IP address of a target's device.
VLAN hopping	An attack in which the source MAC address is changed on frames sent by the attacker.
Double tagging	An attack in which the attacking host adds two VLAN tags instead of one to the header of the frames that it transmits.
MAC spoofing	An attack in which the source MAC address is changed in the header of a frame.
Dynamic Trunking Protocol (DTP)	An unsecure protocol that could allow unauthorized devices to modify a switch's configuration.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	4.1 Given a scenario, apply common security techniques to computing resources.
	• Hardening targets
	  ○ Switches
	4.5 Given a scenario, modify enterprise capabilities to enhance security.
	• Implementation of secure protocols
	  ○ Protocol selection
	  ○ Port selection
	  ○ Transport method

TestOut Security Pro	2.2 Harden Network Devices 2.2.1 Configure and access a switch

5.9.1 Switch Features (Lesson Video)

Transcript:

In this lesson, I'm going to discuss switches and some of their features. This isn't meant to be an exhaustive discussion of switch configurations. Instead, I'll cover topics that can affect your network security. As you probably already know, a switch is a connectivity device that provides communication between different network devices. Switches often use a wired communication medium such as twisted pair and fiber optic to carry the signaling to and from their connected devices. Switches are full-duplex discretionary interconnects that send information only to the intended recipient. I'll explain all of this further in this lesson.

First, let's talk about how a switch works. Every device that connects to a network, whether wired or wireless, has a unique identifier called a Media Access Control address, or MAC address. The MAC address is a 48-bit number separated into six 2-byte numbers. The first 24 bits represent the device manufacturer. In this example, 18-66-DA represents Dell, Inc. The last 24 bits are a random sequential number assigned by the manufacturer to ensure that each device has a unique MAC address. When referencing the OSI model, a MAC address is a layer 2 address and the data that traverses the wire is called a frame. The switch keeps track of the MAC address to which it's connected. Inside every switch is a table of MAC addresses and the port the device is plugged into. This is called the MAC address table. As data traverses the switch, it builds this table based on the frames passing through it. This process continues until all devices connected to the switch are accounted for and the MAC address table is complete. When the MAC address table is complete, the switch can send data directly to the appropriate end device since the table defines which port each device is connected to.
It's also important to know that switches are full-duplex devices. This means they can send and receive data simultaneously. This, along with the data separation by port, means that the chance of an Ethernet collision happening is reduced significantly and speed and efficiency are increased.
Switches are classified as being either unmanaged or managed. An unmanaged switch is often found in homes and small office networks and provides the ability to connect all your devices. Unmanaged switches, on the other hand, aren't configurable. This is just fine for smaller networks that really don't have a need for much sophistication. But larger networks do need additional capabilities. This is where you'll require a managed switch. Managed switches usually have a web-based or command line interface that lets you configure how the switch functions.
For example, a switch's configuration tools let you create and manage virtual LANs, or VLANs. VLANs segment data into different zones. This segmentation might be required for such things as voice, surveillance, and storage. Or it could simply separate zone data for additional resiliency and privacy. Managed switches also support link aggregation, which

provides the ability to combine multiple switch ports in parallel to increase total available bandwidth. This is often used when connecting to devices with multiple NICs that provide things like file and storage services. You can also configure a switch for port security. This lets you define things like which devices and how many can be attached to a single port. A managed switch has many additional benefits that go beyond the scope of this course. But the basic takeaway here is most organizations use managed switches since they can be configured and provide better security.

Another related topic is switch aggregation. Let's look at that a little closer. Often, when you create a switching solution, it's necessary to employ multiple switches throughout a building. There are several reasons for this. The most prevalent is the need to have close proximity to end devices. Ethernet specifications provide for a maximum length of 100 meters between a switch and an endpoint, assuming there are no breaks in the connection such as a wall plate or patch panel. Often, we place wiring closets throughout a building so the 100-meter limit isn't violated.

Normally, an access level switch is placed in the wiring closet to connect a system that's configured for VLANs, port security, and voice over IP. Depending on the number of connections you're dealing with, you might deploy multiple switches for this purpose. These switches are then connected to an upstream switch. The upstream switch is usually much faster and more efficient. Often, the connections between the upstream switch and the access switches are fiber optic, which doesn't have the same distance limitations as a twisted pair of Ethernet cables. In addition, fiber-optic connections are much faster than twisted-pair Ethernet connections.

When you make your decision on the type of switch to purchase, you'll have to choose between a layer 2 and layer 3 switch. While you use both switches to connect end devices to a network, they each have different features.

For example, as the name implies, a layer 2 switch operates at the second layer of the OSI model, which is the Data Link layer. This means the switch understands the connected end devices' physical or MAC addresses. At this level, data passing through switches are known as frames. The port used to send data to an end device is determined by that device's MAC address.

A layer 3 switch operates at the third layer of the OSI model, which is the Network layer. These types of switches can perform the same functions as layer 2 switches. But in addition, a layer 3 switch is aware of the end device's IP address. At the Network layer, the data passing through the switch is known as a packet. A layer 3 switch is also capable of limited IP routing, meaning that layer 3 routers are often used to route packets to and from different VLANs. In practice, Layer 3 switches aren't intended for connecting end devices. Instead, we typically use them for upstream switches.

Now that you understand the basic functions of a switch, let's discuss switch security. Just like any network device, threat actors try to exploit any weakness they can. Switches are no exception. Attackers can use a variety of techniques to compromise switches to gain access to a system.

As such, it's important for system administrators to take precautions and disable switch functions that aren't in use and to monitor switch activity. Three security issues to look out for are MAC flooding, switch spoofing, and ARP poisoning. Let's look at each.

We went over how a switch maintains a table of MAC addresses and the ports to which they're connected. A switch maintains enough memory to store that data and not much more. Knowing this, an attacker might try to flood the switch with bogus MAC addresses. This pushes the legitimate MAC address out and fills up the switch's memory. In effect, this breaks the switch. When the memory is full, the switch reverts to an open state where data is sent to all connected devices. This means that any device connected to the switch is able to capture all data coming out of it. After launching a successful MAC flooding attack, a malicious user can use a packet analyzer to capture sensitive data being transmitted between computers. This wouldn't be possible were the switch operating normally.

Now, let's look at switch spoofing. As we've already discussed, VLANs are used in most organizations to segregate data. Basically, a single network can be broken into smaller chucks using VLAN segmentation. While this is a good practice, it can potentially lead to problems if not done correctly. An attacker could use a less secure segment to access a more secure one. We often configure switches to host multiple VLANs so that end devices that are connected to the right VLAN can communicate with other devices on their individual segment.

Switches eventually do need to communicate with other switches, though. To do this, the switches must be able to transfer information across all VLANs. Someone needs to configure a port as a trunk port, which is then connected to the receiving switch's trunk port. By default, trunk ports have access to all VLANs and pass traffic for multiple VLANs across the same physical link. The attacker takes advantage of this by using Dynamic Trunking Protocol messages, or DTP messages, to spoof itself as a switch. After they accomplish this, the attacker is able to capture all data for all the attached VLANs. One way to prevent a switch spoofing attack is for you to turn off trunking on all ports except the ones that require it. In addition, you should disable DTP on ports that do require trunking, and manually enable the trunking yourself.

The last switch security issue we'll discuss is known as ARP poisoning. Address Resolution Protocol, or ARP, associates a device's MAC address with the IP address assigned to it. To aid in this process, an ARP table is created and stored in a switch's memory for easy lookup. Knowing this, attackers send commands to the switch to overwrite the switch's ARP

table contents by replacing a good, known MAC address with the attacker's device's MAC address. As a result, data that was intended for the legitimate device is now sent to the attacker's device. In many cases, the attacker captures the data and sends it on to its rightful owner. This is known as an on-path attack.
ARP poisoning attacks are difficult to mitigate. If an attacker has access to the internal network, there's little an administrator can do. Specialized training might be a good idea to help you avoid this type of attack. When you work over the internet, you should try to use virtual private networks, or VPNs. VPNs use an encrypted tunnel that largely blocks your activity from ARP spoofing.
That's it for this lesson. This lesson introduced many features we have available on modern switches. We discussed how a switch works, how you can aggregate switches, and the differences between a layer 2 and a layer 3 switch. We ended this lesson by discussing three potential security concerns, which are MAC flooding, switch spoofing, and ARP poisoning.

5.9.2 Securing Network Switches (Lesson Video)

Transcript:

In this lesson, I'll discuss several of the best practices you can implement to help secure your network switches. I'll also cover additional topics that further lock down a switch and add additional controls to make it more secure and less susceptible to an attack. These topics include port security, DHCP snooping, and dynamic ARP inspection.

The first best practice in securing a switch is to change the switch's default credentials. When doing so, choose a username that's not too easy to guess and a complex password using lower and uppercase characters, numbers, and non-alphanumeric symbols.
You should enable secure shell, or SSH, for remote access and also disable telnet. Telnet uses cleartext for usernames and passwords, which isn't good. When it comes to securing a switch, Simple Network Management Protocol, or SNMP, is a powerful ally when you configure it correctly. When it's left unconfigured, though, it presents a security vulnerability that can simply be avoided by disabling it. In the unlikely event there are switch ports left open with no connection, ensure those ports are disabled to prevent unwanted connections.
In addition, you should add a warning banner that states that unauthorized switch configuration access is prohibited. The last best practice is to enable logging. This allows administrators to see any major or minor events a switch experiences. You can also use a logging server to store, archive, and consolidate logging for the enterprise.
Now, let's shift gears to port security. It's up to the switch manufacturer to define how switch security is implemented and the command syntax that's used for setup and configuration. But you can use port security as more granular approach to switch security. Port security configurations can limit the number of devices allowed per port, limit hardware address connections, and define what happens when the port experiences a violation of these settings. For example, you could set it to shut down after a breach. An important note here is that on many switches, port security can't be configured on trunk, EtherChannel, or port analyzer ports.
Since organizations face constant battles with finding open switch connections in a growing workforce, one solution is to add an intermediate switch. In this case, a small switch is placed in a small office that allows several connections, and then that switch is connected to the upstream switch. If the number of devices is limited on the upstream switch and set to one, only the first device that tries to connect is allowed a connection. When the second device tries to connect, the switch denies the connection and the second device won't connect to the network
You can also configure switch ports to only accept connections from specific MAC addresses.
For example, let's configure port 1 for the MAC address ending in BA, and port 2 for the MAC address ending in 45. Assuming the devices connected to those ports match the required MAC address, they're
allowed to connect. While this is a good defense, it's difficult to maintain and very simple to defeat if one device is swapped out for another. The port for that device will need to be reconfigured if a new computer is bought or the network card is changed. Likewise, it's very easy to spoof a MAC address on a given system, rendering MAC address security ineffective.
When a port security rule violation occurs, the switch may react differently depending on its configuration. This might be implemented differently by different switch manufacturers, but the actions are usually closely related. For the least restrictive action, the port simply drops frames from the offending device. The port or interface remains operational and nothing is reported in the logs.

Another setting tells the switch to not only drop the frame but send an alert and log the offending operation. The third, most restrictive violation disables the port when a violation is detected. This requires intervention to re-enable the port and is the default violation mode.

There are several additional measures you can take to protect network switches. Many are common sense precautions, such as physical security and the best practices listed earlier in this lesson. Let's discuss two additional security configurations you might consider.

First, is DHCP snooping. DHCP, as a service, dynamically distributes network parameters, such as IP addresses and default gateways to network devices. For example, when you first start a workstation, it automatically requests and receives the IP information from a DHCP server. In this type of environment, the DHCP server is known as a trusted device, in the same way that routers, firewalls, and file servers are trusted. A trusted device simply means that it's managed by an administrator in your company.

Since most organizations use DHCP for dynamic IP address allocation, attackers often try to place a rogue DHCP server on the network in an attempt to provide end devices with forged IP address information to introduce vulnerabilities onto the network. DHCP snooping attempts to prevent this from occurring. But for DHCP snooping to work, it must be enabled on your devices. While the actual configuration for DHCP snooping is beyond the scope of this lesson, it's important to learn which process to use to enable and configure DHCP snooping in your infrastructure. Know that the configuration process will vary by manufacturer. Once configured, the DHCP snooping process detects, blocks, and drops DHCP packets when the DHCP information comes from an untrusted source. This information is also logged so further actions can be taken to locate and eradicate the rouge DHCP server.

Another security feature available on many switches is called dynamic ARP inspection. This feature rejects invalid or malicious ARP packets, which can help prevent on-path attacks.

An on-path attack is when an attacker secretly relays or alters the communication between two parties who believe they're directly communicating with each other. So how does an attacker implement a on-path attack?

In normal operations, ARP tables are used to identify end devices. Once identified, communications can flow between the devices.

So an attacker assumes the identity of both endpoints and inserts themself in the middle of the communication. Let's see how that works. First, they poison the ARP cache of end device B and tell B that he or she is A. Then they tell A that he or she is B. Once completed, the original connection between A and B is replaced with A sending information to M and M forwarding the information to B, and vice versa. At this point, M sees all the communication traffic between A and B. He or she can log the information or change it depending on the intent. Adding dynamic ARP inspection can help to prevent this type of attack.

To do this, first configure DHCP snooping on your switch as we just discussed. Next, you need to enable dynamic ARP inspection, or DAI. Once DAI is enabled and configured, only trusted sources can affect the ARP table. So now the switch drops the ARP packet when an untrusted source, in this case the attacker, tries to poison the ARP table. That's because the sender's MAC address and sender's IP address don't match an entry in the DHCP snooping database.

That's it for this lesson. In this lesson, we discussed several best practices to ensure switch security. Next, we talked about how port security can be used to prevent unauthorized access to a network. We also discussed how you can use DHCP snooping in conjunction with dynamic ARP inspection to thwart rogue DHCP servers.

5.9.3 Switch Security Facts

This lesson covers the following topics:

- Security switch features
- Implement switch security
- Spanning Tree Protocol

Security Switch Features

The following table lists switch features that can be implemented to increase network security:

Feature	Description

Virtual LAN (VLAN)	A virtual LAN (VLAN) is a logical grouping of computers based on switch ports.
	
	• VLAN membership is configured by assigning a switch port to a VLAN.
	• A switch can have multiple VLANs configured on it, but each switch port can only be a member of a single VLAN (see exception described below).
	• VLANs can be defined on a single switch or configured on multiple interconnected switches. With multiple switches, each switch can be configured with the same VLANs, and devices on one switch can communicate with devices on other switches as long as they are members of the same VLAN.
	• A trunk port is used to connect two switches together.
	  ○ Typically, gigabit ethernet ports are used for trunk ports, although any port can be a trunk port.
	  ○ A trunk port is a member of all VLANs defined on a switch and carries traffic between the switches.
	  ○ When trunking is used, frames that are sent over a trunk port are tagged by the first switch with the VLAN ID so that the receiving switch knows to which VLAN the frame belongs.
	  ○ The Trunking Protocol describes the format that switches use for frame tagging with the VLAN ID.
	  ○ Because end devices do not understand the VLAN tags, the tag is removed from the frame by the switch before the frame is forwarded to the destination device.
	  ○ VLAN tagging is only used for frames that travel between switches on the trunk ports.
	• Using VLANs, the switch can be used to create multiple IP broadcast domains. Each VLAN is in its own broadcast domain, and broadcast traffic is sent only to members of the same VLAN.
	• In a typical configuration with multiple VLANs, workstations in one VLAN can not communicate with workstations in other VLANs. To enable inter-VLAN communication, you need to use a router (or an OSI Layer 3 switch).

MAC filtering/port security	With switch port security, the devices that can connect to a switch through the port are restricted. • Port security uses the MAC address to identify allowed and denied devices. • On the switch, MAC addresses are stored in RAM in a table and are associated with the port. The table can be manually configured, or learning devices can automatically build the table. • You can specify that only a single MAC address is allowed per port, or you can configure each port to allow multiple addresses. • With automatic configuration, the next device or specified number of devices can connect to the port and additional devices are denied. • A port violation occurs when an unauthorized device tries to connect. The switch configuration determines how the switch handles frames from an unauthorized device. The switch can either drop all frames from the unauthorized device or shut down the port, disabling all communications through that port.

Port authentication (802.1x)	Port authentication is provided by the 802.1x protocol and allows only authenticated devices to connect to the LAN through the switch. Authentication uses user names and passwords, smart cards, or other authentication methods.

Implement Switch Security

Be aware of the following when implementing switch security:

- Creating VLANs with switches offers many administrative benefits. For example, you can:
  - Create virtual LANs based on criteria such as workgroup, protocol, or service.
  - Simplify device moves (devices are moved to new VLANs by modifying the port assignment).
  - Control broadcast traffic based on logical criteria (only devices in the same VLAN receive broadcast traffic).
  - Control security (isolate traffic within a VLAN).
- When you use switches to create VLANs, you still need routers to:
  - Route data between VLANs.
  - Provide port filtering. Port filtering filters network packets in and out of devices based on their application type or port number.
  - Route data into and out of the local area network.
- VLANs are commonly used with voice over IP (VoIP) to distinguish voice traffic from data traffic. Traffic on the voice VLAN can be given a higher priority to ensure timely delivery.
- MAC filtering uses the MAC address of a device to drop or forward frames through the switch. Port authentication requires that the user or device authenticates before frames are forwarded through the switch.
- In general, all switch ports are enabled by default. To increase the security of the switch and network, you should disable individual ports that are not in use.

Restricting access by MAC address is difficult to manage and still prone to spoofing. Better security is obtained by forcing computers and/or users to authenticate before full network access is granted. The IEEE 802.1X Port-based Network Access Control (P NAC) standard allows a switch to require authentication when a host connects to one of its ports. 802.1X uses authentication, authorization, and accounting (AAA) architecture:

- Supplicant — the device requesting access such as a user's PC or laptop.
- Authenticator — the switching device (or any type of network access appliance). This does not validate authentication request s directly but acts as a conduit for authentication data.
- Authentication server —the server that holds or can contact a directory of network objects and that can validate authentication requests, issue authorizations, and perform accounting of security events.

The 802.1X standard is implemented by two protocols:

- Extensible Authentication Protocol (EAP) — provides a framework for deploying multiple types of authentication methods. It is often used with digital certificates to establish a trust relationship and create a secure tunnel to transmit the user credential or to perform smart-card authentication without a password.
- Remote Authentication Dial-In User Service (RADIUS) —allows the authenticator and authentication server to communicate authentication and authorization decisions. The authenticator is a RADIUS client; the authentication server is a RADIUS server.

When a host connects to an 802.1X-enabled switch port, the switch opens the port for the EAP over LAN (EAPoL) protocol only. The switch port only allows full data access when the host has been authenticated. The switch receives an EAP packet with the supplicant's credentials. These are encrypted and cannot be read by the switch. The switch uses the RADIUS protocol to send the EAP packet to the authentication server. The authentication server can access the directory of user

accounts and can validate the credential. If authentication is successful, it informs the switch that full network access can be granted.

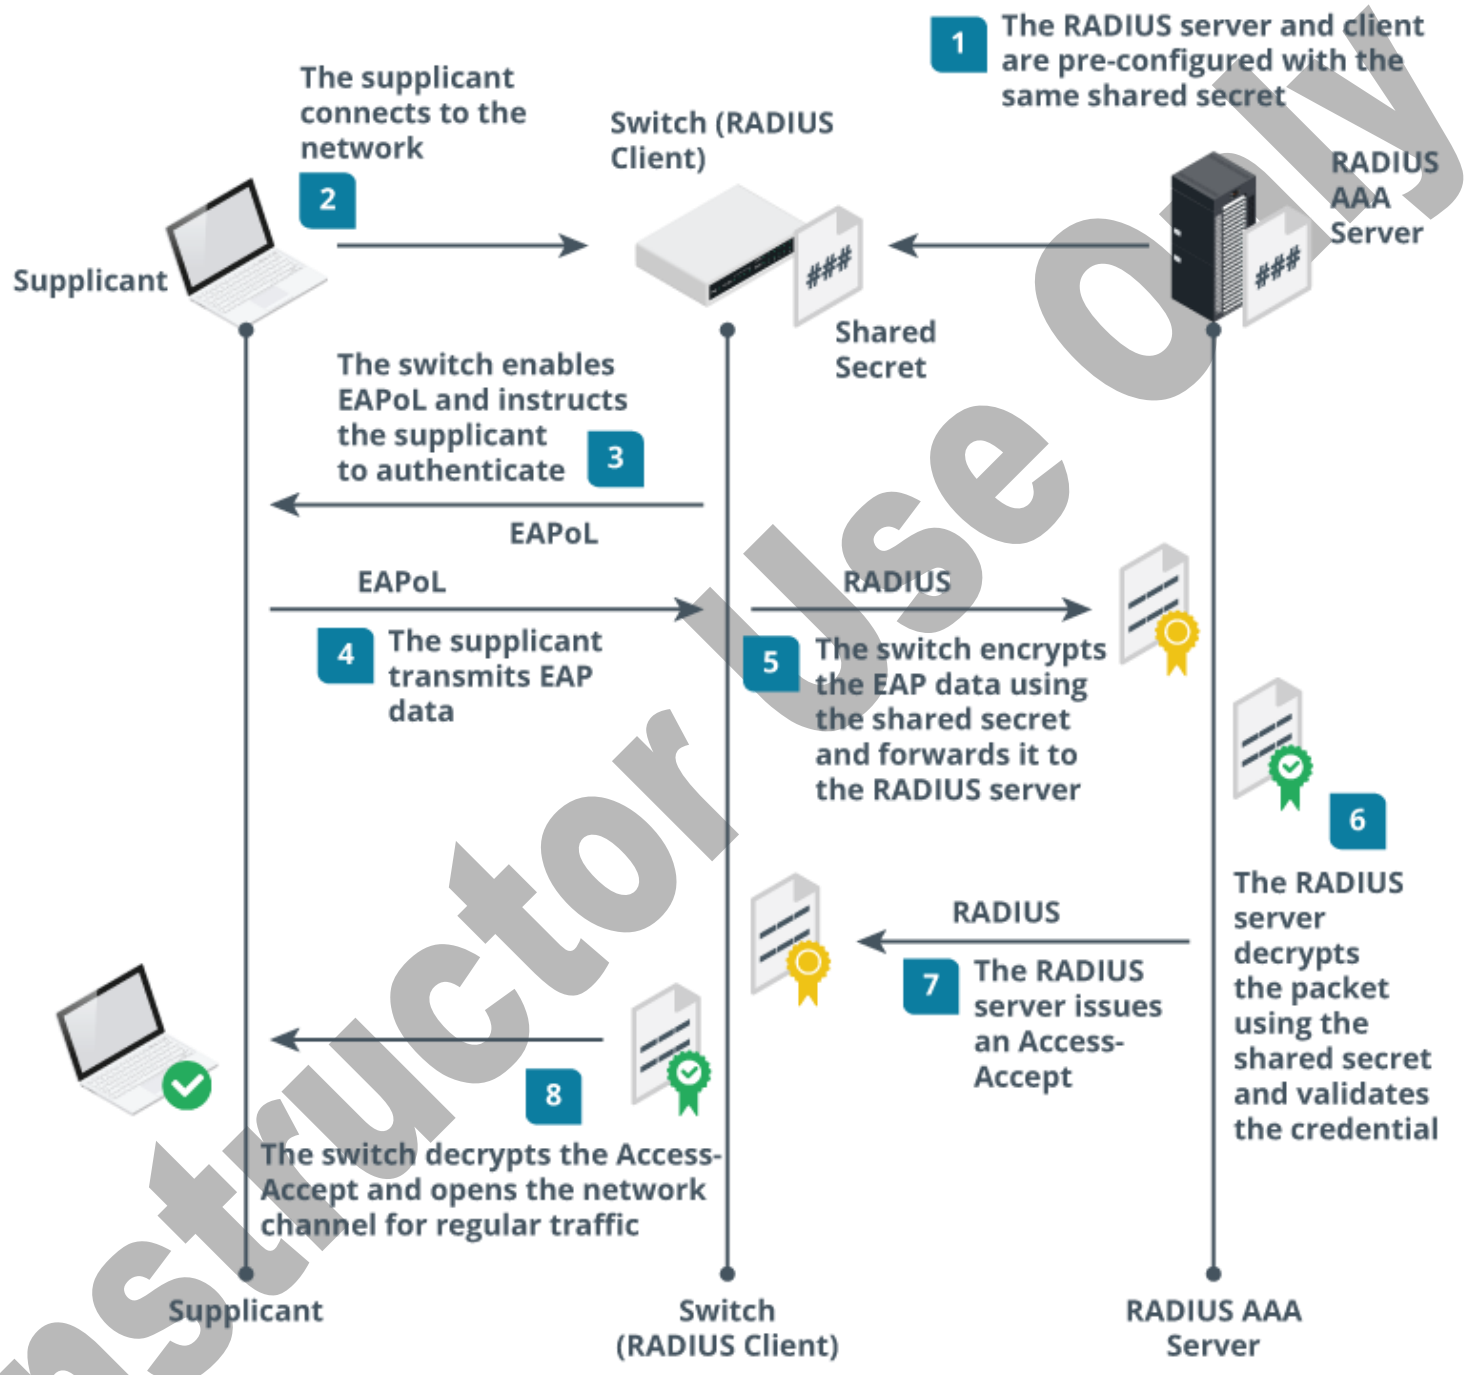

Some hosts are so security-critical that it is unsafe to connect them to any type of network. One example is the root certification authority in PKI. Another example is a host used to analyze malware execution. A host that is not physically connected to any network is said to be air-gapped .

It is also possible to configure an air-gapped network. This means that hosts within the air-gapped network can communicate, but there is no cabled or wireless connection to any other network. Military bases, government sites, and industrial facilities use air-gapped networks.

Physically isolating a host or group of hosts improves security but also incurs significant management challenges. Device administration has to be performed at a local terminal. Any updates or installs have to be performed using USB or optical media. This media is a potential attack vector and must be scanned before allowing its use on an air-gapped host.

Spanning Tree Protocol

To provide fault tolerance, many networks implement redundant paths between devices using multiple switches. However, providing redundant paths between segments causes packets to be passed between the redundant paths endlessly. This condition is known as a switching loop. Switching loops lead to incorrect entries in a MAC address table, making a device appear to be connected to the wrong port and causing unicast traffic to be circulated in a loop between switches. The Spanning Tree Protocol runs on switches to prevent switching loops by making only a single path between switches active at a single time.

The Spanning Tree Protocol also:

- Provides redundant paths between devices.
- Recovers automatically from a topology change or device failure by unblocking redundant paths.
- Identifies the optimal path between any two network devices.
- Calculates the best loop-free path through a network by assigning a role to each bridge or switch and by assigning roles to the ports of each bridge or switch.

The type of ports used by the Spanning Tree Protocol are:

- Root ports, which are configured to communicate directly to the root switch.
- Designated ports, which forward frames to and from attached hosts.
- Blocked ports, which form a loop and are used for redundancy.

Ports in the Spanning Tree Protocol exist in one of five states:

- Blocking
- Listening
- Learning
- Forwarding
- Disabled

5.9.4 Switch Attacks (Lesson Video)

Transcript:

Threat actors look for any opening they can find, and network switches are no exception. Inside attacks take advantage of anything accessible, and attackers can exploit switches to gain valuable information. It's important for you to know that all network information flows through switches and routers, including such things as authentication information and data. Attackers can use a variety of switch attacks methods to gather valuable information that they can use later to exploit the network.

In this lesson, we'll go over the following attack methods: MAC flooding, ARP spoofing, VLAN hopping, STP manipulation, double tagging, and MAC spoofing in regard to Layer 2 switches. Remember, Layer 2 switches understand Ethernet frames at the OSI Data Link layer.
The first attack to examine is MAC flooding. Switches maintain MAC address tables, sometimes called Content Access Memory tables, or CAM tables, to track workstations' and ports' associations. Switches learn which workstation is

connected to each port, and that information is stored in the MAC address table and used when switches forward information.

For example, let's say that Workstation A wants to send information to Workstation D. The frame has a source address ending in As, and the destination address is the address ending with Ds. When A sends the information, the switch examines the destination address and sends the data out port 1 to its destination. Workstations B and C aren't involved, so they never see the information.

Suppose an attacker wants to see all data passing through this switch. Considering that a switch's normal operation only sends data to the intended recipient, the attacker would only be able see data destined for his or her workstation, which is very unlikely. But the attacker can flood the switch with fake MAC address information and fill up the MAC address table. Once the switch can no longer store any more MAC addresses, the switch enters a fail-open mode and behaves the same as a network hub or Layer 1 device. This means that any packet coming in now go out to all switch ports.

The next switch attack is called ARP spoofing, also known as ARP poisoning. But let's first talk about Address Resolution Protocol, or ARP, which works similarly to DNS. DNS resolves a fully qualified domain name into an IP address, while ARP associates a device's MAC address with the IP address assigned to that same device. To aid in this process, a switch creates and stores an ARP table in memory for easy lookup.

Knowing this, an attacker sends commands to the switch to overwrite the contents of its ARP table by replacing a good, known MAC address in the table with the MAC address of the attacker's device. This is the ARP spoofing or ARP poisoning. As a result, data that was intended for the legitimate device is now sent to the attacker's device. In many cases, the attacker captures the data and sends it on to its rightful owner. Another attack method you should know about is VLAN hopping. VLANs are used to separate traffic into sections that function like networks. This is often done to all the different VLANS that carry different types of traffic, such as data, voice, or surveillance information. Since each VLAN is its own network, a device on one VLAN can't communicate with a device on another VLAN without a router. There are a couple of techniques attackers can use to take advantage of this VLAN separation.

The first is VLAN spoofing. VLAN spoofing takes advantage of a switch that's using its default settings, which allows for dynamic trunk negotiation. Depending on the switch, this default setting is known as dynamic auto or dynamic desirable. This means that if you have a switch connected to the original switch, the ports connecting the two are automatically configured as a trunk connection.

A trunk connection allows data from one switch to pass to another. A spoofing attach takes advantage of this by fooling the switch into thinking the attacker is connected using a trunk line. Although there are several ways to carry out a spoofing attach, such as adding a new switch, we'll look at how someone could accomplish a spoofing attach with just one switch.

Our switch here is configured with two VLANs VLAN 10 and VLAN 20 with both having several computers connected to it. Our attacker, a disgruntled employee, is connected to VLAN 10 and wants to steal sensitive information from the company's chief executive officer. But as mentioned, VLAN 10 is a separate network from VLAN 20. So as-is, the attacker can't access the CEO's computer.

By using a tool like Yersinia, our disgruntled employee can send the switch a packet that tells the switch the port he or she is using is really connected to another network switch. This is done using a Dynamic Trunking Protocol frame, or DTP frame. Since the switch is configured to auto-negotiate for a trunk line, it assumes the attacker really is a switch and opens his or her port as a trunk line. With the attacker's port now configured this way, he or she can easily capture packets from the CEO's computer. Keep in mind that the best way to protect your network from this type of attack is to disable the switch's auto-negotiation setting for all ports except those to which an authorized switch is attached.

Another method attackers use is called double tagging. This VLAN hopping technique is used when the attacker's computer is on one VLAN switch, and the target computer is on a VLAN attached to a separate switch. In addition, the switches must be configured to use what's called native VLANs.

When a packet is sent to a switch, it includes what's known as a VLAN tag. This tag's purpose is to indicate which switch should process the frame. For example, you'll notice that both users are on VLAN 10. If User A wants to send data to User B, the frame would include the VLAN tag indicating that the frame was intended for VLAN 10. When the data reaches the switch, the VLAN tag is removed and the packet is sent to its rightful destination.

Double tagging is accomplished by manipulating the frame being sent to include two VLAN tags. So one for VLAN 10 and one for Switch 2, or VLAN 20. When this type of frame is received by VLAN 10's switch, the first VLAN tag is removed. But seeing that the second VLAN tag is referencing VLAN 20, the frame is forwarded to that switch. Then, as it normally would, this switch strips off the second VLAN tag and sends the frame to the victim's VLAN, successfully hopping from one VLAN to another. To prevent double tagging, it's important to make sure that your switches aren't configured to use native VLANs. This is because trunk ports configured with a native VLAN won't apply their own VLAN tag when sending these frames, which lets the attacker's tagged frames to continue as I just described.

Another common switch attack is known as a Spanning Tree Protocol, or STP, manipulation attack. STP is normally configured on a network with several switches. The primary purpose is to prevent switching loops. Often, whether

intentional or not, there are several paths data can take to get from its source to its destination. STP dynamically turns off certain switch ports to ensure that data can't get stuck hopping from switch to switch without ever reaching its intended target.
While STP configuration is beyond the scope of this lesson, it's important to understand that a single switch is designated as the root bridge. The root bridge is an optimized data path's primary source. This is necessary because some switches often connect to other switches to form redundant connections to ensure communications continue even if a switch or port fails. The root bridge is used to pass data from switch to switch since all switches that participate in the tree know where the root bridge is located.
So the root bridge is responsible for calculating the spanning tree from topology changes advertised by non-root bridges. If an attacker can become the root bridge, he or she is then able to see a variety of frames that they normally wouldn't see. To perpetrate this attack, the attacker inserts their switch into the tree and manipulates it to appoint his or her switch as the root bridge. By doing this, he or she can use a sniffer to collect data traversing the network. An attacker accomplishes this manipulation by sending bridge ID frames, or BID frames, with a lower ID than that currently being used by the legitimate root bridge.
There are several step you can take to mitigate this type of attack. First, make sure that attackers can't easily guess which bridge ID number is being used by the legitimate root bridge. Most switches have a default root bridge ID, ensuring that your bridge ID is considerably lower than the default one. If your switch supports it, a second option is to set up Bridge Protocol Data Units Guard, or BPDU Guard. This feature allows an interface to put itself into blocking state when it receives a BPDU packet meant to change the root bridge switch. Your third option is to enable Root Guard on the ports not being used as trunk lines. This keeps ports in their assigned roles. If one of these ports receives a BPDU frame, a sys error is logged and that port is blocked, thwarting the attacker's attempt to change the root bridge.
The last common switch attack to discuss is known as MAC spoofing. MAC spoofing is the process of sending out data from a computer using a MAC address that's different from the MAC address physically hard-coded on a network interface card, or NIC. Although you can't physically change your MAC address, there are tools like Windows Network and Sharing Center in Control Panel, SMAC for Windows, and others that can make an operating system believe that the NIC has a different MAC address. There are several reasons attackers may use this method to attack a network, such as to defeat switch port security.
For example, in some networks, part of a switch's security is the creation of a whitelist of MAC addresses. Only devices with a MAC address in the whitelist can be processed through that switch. If an attacker can use a network sniffer to find a MAC address that's being used by a legitimate host attached to the switch, he or she can change their MAC address to a valid one on the whitelist. Then they've successfully gained access to the switch and any data being transmitted through it.
An attacker could use MAC spoofing to change their MAC address to mimic a targeted system as well. For example, once an attacker has assumed a valid computer's MAC address, he or she could send the switch a MAC update command, changing the MAC address table. Now, the switch sends any data that was destined for the target system to the attacker's system. The important thing to know here is that since it's controlled on the attacker's computer, there really isn't a way to prevent MAC spoofing. If this becomes an issue, you'll need to use techniques at the Application level of the OSI model.
As we end this lesson, you've probably noticed that for any of these attacks to work, the attacker must have physical network access. So physical security is a must! One thing to keep in mind is that for every offense, there's a defense. Administrators can make configuration settings that help to avoid some of these attacks or, at least, notify them something strange is going on.
That's it for this lesson. In this lesson, we examined several methods attackers utilize to exploit switch configuration weaknesses, such as the ones shown here. With each of these attacks, one of the primary goals is to collect information that can be used to exploit vulnerabilities at a later time.

5.9.5 Switch Attack Facts

This lesson covers the following topics:

- Common switch attacks
- Switch hardening
- Implementing secure protocols

Common Switch Attacks

The following table describes common attacks that are perpetrated against switches.

Attack	Description
MAC flooding	*MAC flooding* overloads the switch's MAC forwarding table to make the switch function like a hub. MAC flooding is performed using the following method: 1. The attacker floods the switch with packets, each containing a different source MAC address. 2. The flood of packets fills up the forwarding table and consumes so much of the memory in the switch that it causes the switch to enter a state called *fail open mode* . In fail open mode, all incoming packets are broadcast out to all ports (as with a hub), instead of to the designated port (as a switch normally does). 3. The attacker then captures all the traffic with a protocol analyzer/sniffer.

ARP spoofing/poisoning	*ARP spoofing/poisoning* associates the attacker's MAC address with the IP address of victim's device.
	• When computers send an ARP request for the MAC address of a known IP address, the attacker's system responds with its MAC address.
	• The source device sends frames to the attacker's MAC address instead of the correct device.
	• Switches are indirectly involved in the attack because they do not verify the MAC address/IP address association.
	• A default gateway is a prime target because local traffic goes through a default gateway to get to non-local destinations such as the internet.
	• When the attacker's system MAC address receives packets intended for the default gateway, the attacker can:
	  ○ Forward the packets to the actual default gateway (*passive sniffing*).
	  ○ Modify data in the packets before forwarding it (*man-in-the-middle)* .

MAC spoofing	*MAC spoofing* is changing the source MAC address on frames. The attacker's system sends frames with the spoofed MAC address. The switch reads the source address contained in the frames and associates the MAC address with the port where the attacker is connected. MAC spoofing can be used to: • Bypass 802.1x port-based security. • Bypass wireless MAC filtering. • Hide the identity of the attacker's computer or to impersonate another device on the network. • Impersonate a device on the network to capture frames addressed to that device. • Impersonate a valid device on the network to gain network access. For example, to gain access when the switch is using the MAC address to allow or deny a network connection.

Dynamic Trunking Protocol (DTP)	Switches have the ability to automatically detect trunk ports and negotiate the trunking protocol used between devices. The Dynamic Trunking Protocol is not secure and allows unauthorized devices to possibly modify configuration information. You should disable the DTP services on the switch's end user (access) ports before implementing the switch configuration into the network.

Switch Hardening

Examples of changes designed to improve the security of switches from the default settings include the following:

- Change Default Credentials that are well documented and pose a significant security risk.
- Disable Unnecessary Services and Interfaces on a switch or router. Not every service or interface is needed. For example, services like HTTP or Telnet should be avoided.
- Use Secure Management Protocols such as SSH instead of Telnet or HTTPS instead of HTTP.
- Implement Access Control Lists (ACLs) to restrict access to the switch to only required devices and networks.
- Enable Logging and Monitoring to help identify issues like repeated login failures, configuration changes, and many others.
- Configure Port Security helps limit the devices that can connect to a switch port to prevent unauthorized access.
- Strong Password Policies help reduce the risk of password attacks.
- Physically Secure Equipment like keeping devices in a locked room to prevent unauthorized physical access.

Implementing Secure Protocols

Organizations usually follow formal processes when selecting secure protocols to ensure comprehensive documentation and well-informed decision-making. These processes include assessing risks, reviewing policies, and evaluating the security features of different protocols. Organizations may also consult with technical experts or vendors for recommendations. The outcomes of these processes are documented, which is useful for audits and compliance reviews. Additionally, these process outcomes will typically impact security baselines and configuration management systems.

Selecting protocols, assigning ports, setting transport methods, and other security considerations require careful consideration. The first step requires evaluating the data type used and its sensitivity level. Organizations should select secure protocols like HTTPS, SSH, and SFTP/FTPS for transmitting sensitive or private data. Configuring TCP ports depends on the protocol, as standard ports are associated with specific protocols (HTTP commonly uses port 80, HTTPS uses port 443). While default protocol ports can be changed, doing so may complicate configuration and cause potential accessibility issues.

However, many administrators choose to change standard default ports and a method to obscure them. TCP (Transmission Control Protocol) and UDP (User Datagram Protocol) are two principal transport methods. TCP is connection-oriented and provides reliability, ordering, and error-checking, making it suitable for applications requiring high levels of reliability. UDP is connectionless, making it faster than TCP and more suitable for real-time applications like video streaming, telephony, and gaming, where occasional packet loss is less impactful.

When selecting secure protocols, administrators and analysts must consider suitable encryption levels, authentication methods, existing firewalls or other security equipment, and other factors which may impact the operation of the systems and software they are intended to protect. Ultimately, protocol selection requires an optimum balance among security, maintainability, performance, and cost.

5.9.6 Hardening a Switch (Demo Video)

Transcript:

In this demonstration we'll look at various methods to harden a managed switch. In this case we're using a Cisco Small Business Managed Switch and are logged into the management interface using Internet Explorer. We're authenticated as the Admin User so we can make changes. In this demonstration we'll look at port security, management access, and access control using access control lists, ACLs.

To begin, let's look at port security. When hardening a switch, one of the first things we want to do is shut down any unused ports. Click on Configure Port Settings, which will take us directly to the Port Settings page under Port Management.
On this page we can select one of the unused ports. You'll notice we have several ports here that are either down and unused, or up and used. The ones that are unused are currently down. We want to make sure they're administratively down, meaning they have to be turned on before they can be used. Let's take a look at Port Two. It's down.
Let's scroll to the bottom of the page and click Edit. On this page you'll notice that the Administrative Status of this port is actually up, even though it's operationally down. This means there's nothing plugged into it. To make sure that this port remains down and is not used without permission, we'll set the Administrative Status as Down. Click Apply, Close.
To set all other unused ports as Administratively Down, we'll select Port Two then click Copy Settings. In the Copy Settings window we can list all the ports that we want to be down. In this case, we'll type Three, Six, and Eight through 28. Go ahead and click Apply. All of those ports are now administratively down.
Let's look at Port Three for confirmation. Select Port Three and click Edit. You'll notice on this page that the status Administratively Down. So, the status that we copied from Port Two went to the rest of the ports. Click Close.
By setting the ports that we're not using to Administratively Down, those ports will come up at a down state each time the switch starts up. At the end we'll make sure that we save our Configuration File to the Startup Configuration.
Next, we also want to enable port security. To do that, we'll open the Security Tab and click on Port Security. Under this tab we can configure locking mechanism for each of these ports. We'll click on Port One and select Edit.
On this page you'll notice that we can lock the port, meaning we're going to only allow certain MAC addresses to pass traffic through this port. A Classic Lock means that the first device you connect to that port will be remembered and allowed to pass traffic through this port. You can also select Limited Dynamic Lock, which gives you a variable number from one to 256. Let's change this to four.
Now, if you connect a hub to a specific port on the switch and connect four devices to that, you can allow all four of those devices to talk through that port. No additional devices would be allowed to pass traffic.
We'll go ahead and select the Classic Lock. That will give us one device, and the device will be locked in. If you try to connect a device that is not supposed to be on that port, the Action on Violation determines what happens. The default is to discard the packet. If you connect any other device to the port, the traffic will be ignored.
You can also forward the traffic, or you can shut down the port. That's a little drastic for our purposes. However, if you're extremely cautious about your security and you don't want anybody connecting a device you don't know about, being able to shut down the port is one of the most secure things that you can do.
We'll go ahead and select Discard if it's a device that we aren't sure about. If somebody plugs in a device that's not locked in you can also trap the MAC addresses of the offending device. We've enabled locking for Port One, as well as the classic lock. We're going to discard the frames of any other devices that come onto that port. We'll go ahead and click Apply, Close.
Once again, we can take the settings that we've created for Port One and apply them to the rest of the ports. Scroll down to the bottom and select Copy Settings. In this case, we'll select all of the ports, or Two through 28, and click Apply.
You'll see that the settings that we've created for port security now apply to all 28 ports.
The next piece that we would like to cover is management access. We'll go to Management Access Method and click Access Profiles. In this case, the only profile that's on the machine is the Console Profile, which is the default.
We're going to add a new access profile, so click Add. We'll call this access profile Management. In this Management Profile the very first rule, or Rule One that we want to have will deny all management. We'll go ahead and select Deny and specify all different management methods, and apply it to all interfaces. That will be our first and main rule. We'll go ahead and click Apply, Close.
Now that we've created our Management Profile, we can select Profile Rules. It can seem a little confusing that you create a Management Profile then add a rule to it. We'll go ahead and select our Management Profile, and add another rule. It looks very similar to the page we just used.

For Rule Number Two, we want to allow HTTP access. We want to select all interfaces, but we also only want to allow traffic from a specific workstation or IP address. We'll specify that 192.168.1.200 will be the only workstation allowed to access this port. We'll also input the network mask by typing 255.255.255.0. Alternatively, we can use the prefix length and use the prefix instead. We'll go ahead and click Apply, Close.
We now have a management access rule that denies all traffic, and one access rule which allows us to access the HTTP interface from one specific workstation. But, we're not done yet.
We need to go back to Access Profiles and set the Active Profile. Right now, we don't have an active management profile. Go ahead and select Management. This denies all types of management interfaces with this switch, except for the HTTP access from the one workstation that we have selected.
We're not going to click Apply this time. If we do, we'll probably be locked out of the switch, and we want to make sure that we have access to complete this demo.
Let's go to the next section, and look at firmware updates. Under Administration, File Management we can select Upgrade, Backup Firmware. Part of hardening the device or a computer of any kind is to make sure you have the latest software and patches. You can use browse to select the latest firmware and download it.
In this particular switch, you have to select the Active Image once you've loaded the firmware. You can select Image One or the new image, Image Two, and apply it. Next you'll need to reboot the switch. Once you do that the switch will come back with the latest and greatest software and the best enhancements. This allows you to have the best possible security for this switch. Upgrading firmware for the switch as well as applying patches to any computer is a key element to hardening the switch.
The last aspect of hardening the switch is access control using the ACLs. Let's select Access Control. On this switch, we have various methods of access control. We can use a MAC based ACL. We can use an IPV4 based ACL. And we can also use an IPV6 based ACL. We'll set up a MAC based ACL.
In this case, we don't have an access control list, so we'll go ahead and create one. We want to make sure that nobody can attach game consoles or connect any game to our switch. Type Games as the name of our ACL. Click Apply, Close.
We have the ACL, but to put access control entries, ACE, into this access control list, we need to click the MAC Based ACE Table. It has a slightly different interface, but let's click Add to add an entry to this page.
For our first access control entry, Entry One, we're going to select Deny. We'll select Any for the Destination MAC address. Under Source MAC Address, we're going to select User Defined and put in the address, one of the game consoles, or the prefix of the MAC address, which is used for all devices of the same manufacturer. We'll go ahead and enter 00041F followed by six ones. Under Source MAC Wildcard Mask we'll enter six zeroes followed by six ones.
This says that we only care about the first six digits of the MAC Address. This is preventing any device from the same manufacturer from being connected. We'll go ahead and click Apply.
That creates our first entry into the Access Control List Entry Table. We'll add one more just for demonstration purposes. We'll go ahead and create Priority Two.
Once again, we're going to select Deny, and specify the Destination MAC Address as Any. We're going to specify the Source MAC Address value as 005042 followed by six ones. We'll do the same Source MAC Wildcard Mask, or six zeroes followed by six ones. Click Apply, Close to create our access control entries.
We could add more for any of the game console types that are out there, or any other device that we want to prohibit from using our network.
We can also add anything we want to allow on the network. Now that we've created the MAC Based ACL Table, and have created the entries in that table, we need to make sure that the ports have a binding to that table. We'll go ahead and click on ACL Binding. Select Port One, scroll down, and click Edit. We want to apply this MAC based ACL, so we'll select that and then select Games, which we just created. If there were additional ones, we could choose from them as well. We'll go ahead and click Apply, Close.
Just as we've done in the previous portions of this demonstration, we'll select Port One. The GE One interface, and click on Copy Settings. We'll enter two through 28 and click Apply. That gave us a successful binding of the Games ACL to all of the ports in our interface.
One final thing that you will always want to do when you've made configuration changes to your switch is to save the Running Configuration to the Startup Configuration. On this switch, simply click Save at the top. It will blink when there have been changes. Select Running Configuration as the Source File Name. And set the Destination File Name as Startup Configuration, then click Apply.
In this demonstration we've examined the idea of port security. We've looked at management access to our switch. And we've also looked at restricting certain traffic to the switch using ACLs.

5.9.7 Harden a Switch (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You need to increase the security on the switch in the networking closet.

The following table lists the used and unused ports:

Unused Ports	Used Ports
GE2 GE7 GE9-GE20 GE25 GE27-GE28	GE1 GE3-GE6 GE8 GE21-GE24 GE26

In this lab, your task is to:

- Shut down the unused ports.
- Configure the following Port Security settings for the used ports:
  - Interface Status: **Lock**
  - Learning Mode: **Classic Lock**
  - Action on Violation: **Discard**

Explanation

While completing this lab, use the following information:

Unused Ports	Used Ports
GE2 GE7 GE9-GE20 GE25 GE27-GE28	GE1 GE3-GE6 GE8 GE21-GE24 GE26

Complete this lab as follows:

1. Shut down the unused ports.
   a. Under Initial Setup, select **Configure Port Settings** .
   b. Select the **GE2** port.
   c. Scroll down and select **Edit** .
   d. Under Administrative Status, select **Down** .
   e. Scroll down and select **Apply** .
   f. Select **Close** .
   g. With the GE2 port selected, scroll down and select **Copy Settings** .
   h. In the Copy configuration field, enter the remaining **unused ports** . (View the example for the proper syntax.)
   i. Select **Apply** . From the Port Setting Table, in the Port Status column, you can see that all the ports are down now.

2. Configure the Port Security settings.
    a. From the left menu, expand **Security** .
    b. Select **Port Security** .
    c. Select the **GE1** port.
    d. Scroll down and select **Edit** .
    e. For Interface Status, select **Lock** .
    f. For Learning Mode, make sure **Classic Lock** is selected.
    g. For Action on Violation, make sure **Discard** is selected.
    h. Select **Apply** .
    i. Select **Close** .
    j. Scroll down and select **Copy Settings** .
    k. Enter the remaining **used ports** . (View the example for the proper syntax.)
    l. Select **Apply** .

5.9.8 Secure Access to a Switch (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You need to increase the security on the switch in the Networking Closet by restricting access management.

In this lab, your task is to:

- Create an access profile named *MgtAccess* and configure it with the following settings:

Setting	Value
Access Profile Name	MgtAccess
Rule Priority	1
Management Method	All
Action	Deny
Applies to Interface	All
Applies to Source IP address	All

- Add a profile rule to the *MgtAccess* profile with the following settings:

Setting	Value

Rule Priority	2
Management Method	HTTP
Action	Permit
Applies to interface	All
Applies to Source IP address	User defined IP Version: Version 4 IP Address: 192.168.0.10 Network Mask: 255.255.255.0

- Set the *MgtAccess* profile as the active access profile.
- Save the changes to the switch's startup configuration file using the default settings.

If you need to log back into the switch, the username is **ITSwitchAdmin** , and the password is **Admin$only** .

Explanation

Complete this lab as follows:

1. Create and configure an Access Profile named *MgtAccess* .
    a. From the left pane, expand and select **Security** > **Mgmt Access Method** > **Access Profiles** .
    b. Select **Add** .
    c. Enter the Access Profile Name of **MgtAccess** .
    d. Enter the Rule Priority of **1** .
    e. For Action, select **Deny** .
    f. Select **Apply** and then select **Close** .
2. Add a profile rule to the *MgtAccess* profile.
    a. From the left pane, under **Security** > **Mgmt Access Method** , select **Profile Rules** .
    b. From the right pane, select the **MgtAccess** profile and then select **Add** .
    c. Enter a Rule Priority of **2** .
    d. For Management Method, select **HTTP** .
    e. For *Applies to Source IP Address* , select **User Defined** .
    f. For IP Address, enter **192.168.0.10** .
    g. For Mask, enter a Network Mask of **255.255.255.0** .
    h. Select **Apply** and then select **Close** .
3. Set the MgtAccess profile as the active access profile.
    a. From the left pane, under **Security** > **Mgmt Access Method** , select **Access Profiles** .
    b. Use the Active Access Profile drop-down list to select **MgtAccess** .
    c. Select **Apply** .
    d. Select **OK** .
4. Save the changes to the switch's startup configuration file.
    a. At the top, select **Save** .
    b. On the right, under *Source File Name* , make sure **Running configuration** is selected.
    c. For the Destination File Name, make sure **Startup configuration** is selected.
    d. Select **Apply** .

e. Select **OK** .

5.9.9 Secure Access to a Switch 2 (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You need to increase the security on the switch in the Networking Closet by creating an access control list. You have been asked to prevent video game consoles from connecting to the switch.

In this lab, your task is to:

- Create a MAC-based ACL named **GameConsoles** .
- Configure the **GameConsoles** MAC-based access control entry (ACE) settings as follows:

Priority	Action	Destination MAC Address	Source MAC Address
1	Deny	Any	Value: **00041F111111** Mask: **000000111111**
2	Deny	Any	Value: **005042111111** Mask: **000000111111**
3	Deny	Any	Value: **000D3A111111** Mask: **000000111111**
4	Deny	Any	Value: **001315111111** Mask: **000000111111**
5	Deny	Any	Value: **0009BF111111** Mask: **000000111111**
6	Deny	Any	Value: **00125A111111** Mask: **000000111111**

- Bind the **GameConsoles** ACL to all of the **GE1-GE30** interfaces.

  Use **Copy Settings** to apply the binding to multiple interfaces.

- Save the changes to the switch's startup configuration file. Use the default settings.

Explanation

While completing this lab, use the following information:

- Configure the *GameConsoles* MAC-based access control entry (ACE) settings as follows:
-

Priority	Action	Destination MAC Address	Source MAC Address
1	Deny	Any	Value: **00041F111111** Mask: **000000111111**
2	Deny	Any	Value: **005042111111** Mask: **000000111111**
3	Deny	Any	Value: **000D3A111111** Mask: **000000111111**
4	Deny	Any	Value: **001315111111** Mask: **000000111111**
5	Deny	Any	Value: **0009BF111111** Mask: **000000111111**
6	Deny	Any	Value: **00125A111111** Mask: **000000111111**

Complete this lab as follows:

1. Create the GameConsoles ACL.
    a. From the Getting Started page, under Quick Access, select **Create MAC-Based ACL** .
    b. Select **Add** .
    c. In the *ACL Name* field, enter **GameConsoles** .
    d. Select **Apply** and then select **Close** .
2. Create a MAC-based access control.
    a. Select **MAC-Based ACE** .
    b. Select **Add** .
    c. Enter the ***priority*** .
    d. Select the ***action*** .
    e. For Destination MAC Address, make sure **Any** is selected.
    f. For Source MAC Address, select **User Defined** .
    g. Enter the ***source MAC address value*** .
    h. Enter the ***source MAC address mask*** .
    i. Select **Apply** .
    j. Repeat steps 2c–2i for the remaining ACE entries.
    k. Select **Close** .
3. Bind the GameConsoles ACL to all of the interfaces.
    a. From the left pane, under Access Control, select **ACL Binding (Port)** .
    b. Select **GE1** .
    c. At the bottom of the window, select **Edit** .
    d. Select **Select MAC-Based ACL** .

e. Select **Apply** and then select **Close** .
f. Select **Copy Settings** .
g. In the Copy configuration's *to* field, enter **2-30** .
h. Select **Apply** .

4. Save the Configuration.
    a. From the top of the window, select **Save** .
    b. On the right, under *Source File Name* , make sure **Running configuration** is selected.
    c. Under *Destination File Name* , make sure **Startup configuration** is selected.
    d. Select **Apply** .
    e. Select **OK** .

5.9.10 Practice Questions (Section Quiz)

q_sec_swi_802x_01_secp8

You want to increase the security of your network by allowing only authenticated users to access network devices through a switch.

Which of the following should you implement?

Answers:

- ***802.1x authentication**
- Port security
- Spanning Tree Protocol
- IPsec

Explanation:

802.1x authentication is an authentication method used on a LAN to allow or deny access based on a port or connection to the network. 802.1x is used for port authentication on switches and to wireless access points and requires an authentication server for validating user credentials. This server is typically a RADIUS server. Authenticated users are allowed full access to the network, while unauthenticated users only have access to the RADIUS server.

Port security uses the MAC address to allow or deny connections based on the MAC address of the device, not user authentication.

Spanning Tree is a protocol for identifying multiple paths through a switched network.

IPsec is a tunneling protocol that adds encryption to packets.

q_sec_swi_802x_02_secp8

Which of the following scenarios would typically utilize 802.1x authentication?

Answers:

- ***Controlling access through a switch**
- Controlling access through a router
- Authenticating remote access clients
- Authenticating VPN users through the internet

Explanation:

802.1x authentication is an authentication method used on a LAN to allow or deny access based on a port or connection to the network. 802.1x is used for port authentication on switches and requires an authentication server for validating user credentials. This server is typically a RADIUS server.

Remote access authentication is handled by remote access servers or a combination of remote access servers and a RADIUS server for centralized authentication.

VPN connections can be controlled by remote access servers or by a special device called a VPN concentrator.

q_sec_swi_mac_filtering_secp8

As a network administrator, you have implemented MAC filtering as a security measure on your company's network. You notice that an unauthorized device has been able to connect to the network despite the MAC filtering.

Which of the following is the MOST likely explanation for this occurrence and what should be your next step?

Answers:

- The MAC address of the unauthorized device has been manually added to the allowed list. You should review the allowed list and remove any unfamiliar addresses.
- ***The unauthorized device has spoofed the MAC address of an authorized device. You should implement a secondary authentication method to increase security.**
- The MAC filtering feature is malfunctioning. You should troubleshoot the feature or consider replacing the switch.
- The unauthorized device has bypassed the MAC filtering by connecting to a disabled port. You should ensure all unused ports are properly disabled.

Explanation:

The unauthorized device has spoofed the MAC address of an authorized device is the most likely explanation. You should implement a secondary authentication method to increase security. MAC spoofing is a technique where an unauthorized device mimics the MAC address of an authorized device, thereby bypassing the MAC filtering. To prevent this, a secondary authentication method, such as 802.1x port-based authentication, can be implemented.

The MAC address of the unauthorized device has been manually added to the allowed list could be a possibility if someone with access to the MAC filtering settings has mistakenly or intentionally added the unauthorized device's MAC address to the allowed list. However, this is less likely than option B, as MAC spoofing is a common method used to bypass MAC filtering.

While it's possible that the MAC filtering feature could be malfunctioning, this is less likely than the device spoofing a MAC address. Troubleshooting the feature or replacing the switch should only be considered after other more likely possibilities have been ruled out.

If a port is disabled, it should not allow any device to connect through it, so this is not a likely explanation. However, it's always a good practice to ensure unused ports are properly disabled to prevent unauthorized access.

q_sec_swi_port_sec_01_secp8

In which of the following situations would you use port security?

Answers:

- You want to prevent sniffing attacks on the network.
- ***You want to restrict the devices that could connect through a switch port.**
- You want to control the packets sent and received by a router.
- You want to prevent MAC address spoofing.

Explanation:

Use port security on a switch to restrict the devices that can connect to the switch. Port security uses the MAC address to identify allowed and denied devices. When an incoming frame is received, the switch examines the source MAC address to decide whether to forward or drop the frame.

Port security cannot prevent sniffing or MAC address spoofing attacks.

Use an access control list on a router to control sent and received packets.

q_sec_swi_port_sec_02_secp8

You are the network administrator for a city library. Throughout the library are several groups of computers that provide public access to the internet. Supervision of these computers has been difficult. You've had problems with patrons bringing personal laptops into the library and disconnecting the network cables from the library computers to connect their laptops to the internet.

The library computers are in groups of four. Each group of four computers is connected to a hub that is connected to the library network through an access port on a switch. You want to restrict access to the network so that only library computers are permitted connectivity to the internet.

What can you do?

Answers:

- ***Configure port security on the switch.**
- Remove the hub and place each library computer on its own access port.
- Create a VLAN for each group of four computers.
- Create static MAC addresses for each computer and associate each address with a VLAN.

Explanation:

Configuring port security on the switch can restrict access so that only specific MAC addresses can connect to the configured switch port. This would prevent the laptop computers from being permitted connectivity.

Placing each library computer on its own access port would have no effect.

VLANs are used to group broadcast traffic and do not restrict connectivity of devices as needed in this scenario.

q_sec_swi_spanning_secp8

You manage a single subnet with three switches. They are connected to provide redundant paths between the switches.

Which feature prevents switching loops and ensures there is only a single active path between any two switches?

Answers:

- Trunking
- ***Spanning Tree Protocol**
- PoE
- 802.1x
- Bonding

Explanation:

Spanning Tree Protocol is a protocol on a switch that allows the switch to maintain multiple paths between switches within a subnet. Spanning Tree Protocol runs on each switch and is used to select a single path between any two switches.

- Without Spanning Tree Protocol, switches that are connected together with multiple links would form a switching loop where frames are passed back and forth continuously.
- Spanning Tree Protocol provides only a single active path between switches. Switch ports that are part of that path are placed in a forwarding state.
- Switch ports that are part of redundant but unused paths are placed in a blocking (non-forwarding) state.
- When an active path goes down, Spanning Tree Protocol automatically recovers and activates the backup ports necessary to provide continued connection between devices.

Bonding does the opposite of Spanning Tree Protocol. Bonding allows multiple switch ports to be used at the same time to reach a specific destination.

802.1x is an authentication protocol used with port security (or port authentication).

Power over Ethernet (PoE) supplies power to end devices through the RJ-45 Ethernet switch port.

Trunking identifies ports that are used to carry VLAN traffic between switches. A trunk port is a member of all VLANs defined on all switches.

q_sec_swi_vlan_01_secp8

When configuring VLANs on a switch, which type of switch ports are members of all VLANs defined on the switch?

Answers:

- ***Trunk ports**
- Uplink ports
- Any port not assigned to a VLAN
- Gigabit and higher Ethernet ports
- Each port can only be a member of a single VLAN

Explanation:

A trunk port is a member of all VLANs defined on a switch and carries traffic between the switches. When trunking is used, frames that are sent over a trunk port are tagged by the first switch with the VLAN ID so that the receiving switch knows to which VLAN the frame belongs.

Typically, uplink ports (which are faster than the other switch ports) are used for trunk ports (not VLANs), although any port can be designated as a trunk port.

On an unconfigured switch, ports are members of a default VLAN (often designated VLAN 1). When you remove the VLAN membership of a port, it is reassigned back to the default VLAN. Therefore, the port is always a member of at least one VLAN.

The hardware used (gigabit and higher Ethernet ports) is not a factor in assigning switch port to VLANs.

Once again, by default, the port is always a member of at least one VLAN, but can be a member of several VLANs.

q_sec_swi_vlan_02_secp8

Which of the following BEST describes the concept of a virtual LAN?

Answers:

- Devices connected by a transmission medium other than a cable (microwave, radio transmissions).
- Devices in separate networks (different network addresses) logically grouped as if they were in the same network.
- ***Devices on the same network logically grouped as if they were on separate networks.**
- Devices connected through the internet that can communicate without using a network address.
- Devices on different networks that can receive multicast packets.

Explanation:

A virtual LAN is created by identifying a subset of devices on the same network logically grouped as if they were on separate networks. Think of VLANs as subdivisions of a LAN.

Devices connected by a transmission medium other than a cable normally rely on a Wi-Fi connection, not a VLAN.

The devices on a VLAN need to be on the same network--not separate networks.

The devices on a VLAN need to be assigned a network address in order to communicate with each other.

Once again, devices on a VLAN need to be on the same network, whether they can receive multicast packets or not.

q_sec_swi_vlan_03_secp8

The IT department of a medium-sized enterprise is reviewing its network architecture with a focus on increasing security and efficiency. The organization currently uses a flat network model, but the security team has proposed implementing Virtual Local Area Networks (VLANs) to compartmentalize traffic and minimize potential attack surfaces.

The team's goal is to limit lateral movement between network devices and enforce a principle of least privilege across the network.

Considering this security improvement initiative, what is a major benefit of integrating VLANs into the existing network architecture from a security standpoint?

Answers:

- Enhancing bandwidth efficiency and speed
- ***Isolating network traffic and reducing the potential attack surface**
- Improving scalability by adding more devices to the network
- Providing an alternative for physical cabling and switches

Explanation:

VLANs isolate network traffic and reduce the potential attack surface, enabling the logical segmentation of networks into isolated broadcast domains. As a result, they limit lateral movement and reduce the potential attack surface.

While VLANs can contribute to better network management, the main security benefit does not relate to bandwidth efficiency or speed.

Although VLANs can handle more devices compared to a flat network, the question focuses on the security benefits, not network expansion.

While it is true that VLANs can reduce the need for physical cabling and switches, this is more of a cost and management benefit, not a primary security advantage.

q_sec_swi_vlan_04_secp8

A newly established organization has decided to implement Virtual LANs (VLANs) for segmenting workstation computer hosts from Voice over Internet Protocol (VoIP) handsets.

The organization is using two VLANs that map to two subnets: 10.1.32.0/24 for workstation computers and 10.1.40.0/24 for VoIP handsets.

In this setup, what could be a potential security advantage?

Answers:

- ***Enhanced control over communication between VLANs.**
- Unrestricted communication between devices on different VLANs.
- Immunity from unauthorized communication interception.
- Physical security of network devices and connected computers.

Explanation:

One security advantage of implementing VLANs is the ability to apply access control rules that prevent or permit certain types of communication between VLANs, hence mitigating risks.

One of the main purposes of VLANs is to segment the network, meaning devices from different VLANs cannot freely communicate with each other without routing.

While VLANs can increase network security by segregating traffic, they do not make it impossible for unauthorized individuals to intercept communications. The organization would need other security measures, like encryption, for that purpose.

While VLANs contribute to network security, they do not ensure the physical security of network devices and connected computers. The organization would require physical security measures for that purpose.

q_sec_swi_vlan_05_secp8

A network engineer is segmenting a company's network to improve security. In terms of routing infrastructure, which of the following strategies would the engineer employ to segment different types of hosts attached to the same switch?

Answers:

- ***Assign each host to a different virtual local area network (VLAN).**
- Assign a different internet protocol (IP) address to each host on the switch.

- Use the Address Resolution Protocol (ARP) to map each host's IP interface to a different media access control (MAC) address.
- Use the same VLAN IDs and subnets for all hosts.

Explanation:

The network engineer can create separate broadcast domains by assigning each host to a different VLAN. The engineer should route traffic sent from one VLAN to another, providing an additional layer of security and segmentation.

Even though the network engineer assigns a unique IP address to each host, it does not necessarily provide segmentation at the level of the switch.

Although the network engineer uses ARP to map IP interfaces to MAC addresses, it does not serve as a strategy for network segmentation.

The network engineer does not create segmentation within the switch using the same VLAN IDs and subnets for all hosts. The engineer should assign different hosts to different VLANs for effective segmentation.

q_sec_swi_vlan_switches_secp8

As a network administrator, you are tasked with creating VLANs on a switch to improve network performance and security.

You decide to create VLANs based on different departments in your company. However, you notice that devices from different VLANs are unable to communicate with each other.

Which of the following is the MOST likely reason for this and what should be your next step?

Answers:

- The switch is not properly configured for VLANs. You should review the VLAN configuration on the switch.
- The devices are not properly configured for VLANs. You should review the network settings on each device.
- ***Inter-VLAN routing is not enabled. You should implement a router or a Layer 3 switch to enable communication between VLANs.**
- The switch ports are not correctly assigned to the VLANs. You should review the VLAN assignments for each port.

Explanation:

Inter-VLAN routing is not enabled is is the most likely explanation. By default, devices in different VLANs cannot communicate with each other. To enable communication between VLANs, you need to implement a router or a Layer 3 switch for inter-VLAN routing.

The switch is not properly configured for VLANs could be a possibility if the switch is not correctly configured for VLANs. However, if the devices are correctly assigned to their respective VLANs and are able to communicate within their VLAN, it's likely that the switch is correctly configured for VLANs.

The devices do not need to be specifically configured for VLANs. The switch handles the VLAN assignments based on the switch ports, so this is not a likely explanation.

If the switch ports were not correctly assigned to the VLANs, the devices would not be able to communicate even within their VLAN. Since the devices are able to communicate within their VLAN, it's likely that the switch ports are correctly assigned to the VLANs.

q_swi_attack_arp_poison_secp8

Which of the following switch attacks associates the attacker's MAC address with the IP address of the victim's devices?

Answers:

- MAC spoofing
- ***ARP spoofing/poisoning**
- DNS poisoning
- Cross-site scripting (XSS)

Explanation:

ARP spoofing/poisoning associates the attacker's MAC address with the IP address of the victim.

MAC spoofing is changing the source MAC address on frames sent by the attacker.

DNS poisoning occurs when a name server receives malicious or misleading data that incorrectly maps host names and IP addresses.

Cross-site scripting (XSS) attacks are a type of injection attack where malicious code is saved onto an otherwise benign site.

q_swi_attack_harden_switches_secp8

You are a network security analyst for a large corporation. The company has recently experienced a series of network attacks, and you've been tasked with hardening the network switches to prevent future attacks. You've identified several potential measures to improve security.

Which of the following would be the MOST effective approach to hardening the switches?

Answers:

- Changing the default credentials.
- Implementing Access Control Lists (ACLs).
- Enabling HTTP services for remote management.
- ***Disabling unnecessary services and interfaces.**

Explanation:

Disabling unnecessary services and interfaces is the most effective approach. Disabling unnecessary services and interfaces reduces the attack surface of the switch, making it harder for attackers to gain access. This is a more comprehensive measure compared to the other options, making it the most effective approach to switch hardening.

Changing default credentials is a basic step in hardening any device, as default credentials are often well-known and pose a significant security risk. However, this measure alone does not provide a comprehensive approach to switch hardening.

ACLs restrict access to the switch, providing a layer of security. However, this measure alone does not provide a comprehensive approach to switch hardening.

Enabling HTTP services for remote management can actually introduce security vulnerabilities, as HTTP is not a secure protocol. This would not be an effective approach to hardening the switches and is therefore not the most effective approach.

q_swi_attack_l2_attack_secp8

Drag each description on the left to the appropriate switch attack type on the right.

Answers:

- Causes packets to fill up the forwarding table and consumes so much of the switch's memory that it enters a state called Fail Open Mode.
- The source device sends frames to the attacker's MAC address instead of to the correct device.
- Can be used to hide the identity of the attacker's computer or impersonate another device on the network.
- Should be disabled on the switch's end user (access) ports before implementing the switch configuration into the network.

Explanation:

Common attacks that are perpetrated against switches are MAC flooding, ARP spoofing/poisoning, and MAC spoofing.

MAC flooding overloads the switch's MAC forwarding table to make the switch function like a hub.

- The attacker floods the switch with packets, each containing a different source MAC address.
- The flood of packets fills up the forwarding table and consumes so much of the memory in the switch that it causes the switch to enter into fail open mode. While in this mode, all incoming packets are broadcast out of all ports (as with a hub) instead of just to the correct ports, as per normal operation.
- The attacker captures all the traffic with a protocol analyzer/sniffer.

ARP spoofing/poisoning associates the attacker's MAC address with the IP address of victim devices.

- When computers send an ARP request for the MAC address of a known IP address, the attacker's system responds with its own MAC address.
- The source device sends frames to the attacker's MAC address instead of to the correct device.
- Switches are indirectly involved in the attack because they do not verify the MAC address/IP address association.

MAC spoofing changes the source MAC address on frames sent by the attacker.

- MAC spoofing is typically used to bypass 802.1x port-based security.
- MAC spoofing can be used to bypass wireless MAC filtering.
- MAC spoofing can be used to hide the identity of the attacker's computer or to impersonate another device on the network.

Dynamic Trunking Protocol (DTP) switches have the ability to automatically detect trunk ports and negotiate the trunking protocol used between devices. DTP is not secure and allows unauthorized devices to possibly modify configuration information. You should disable the DTP services on the switch's end user (access) ports before implementing the switch configuration into the network.

q_swi_attack_mac_flood_secp8

Which of the following attacks, if successful, causes a switch to function like a hub?

Answers:

- ARP poisoning

- ***MAC flooding**
- MAC spoofing
- Replay attack

Explanation:

MAC flooding overloads the switch's MAC forwarding table to make the switch function like a hub. The attacker floods the switch with packets, each containing different source MAC addresses. The flood of packets fills up the forwarding table and consumes so much of the memory in the switch that it causes the switch to enter a state called fail open mode. While in this mode, all incoming packets are broadcast out of all ports (as with a hub), instead of just to the correct ports, as per normal operation.

ARP poisoning associates the attacker's MAC address with the IP address of victim devices. When computers send an ARP request to get the MAC address of a known IP address, the attacker's system responds with its own MAC address. MAC spoofing is changing the source MAC address on frames sent by the attacker.

In a replay attack, the attacker uses a protocol analyzer or sniffer to capture authentication information going from the client to the server. The attacker then uses this information to connect at a later time and pretend to be the client.

q_swi_attack_mac_spoof_secp8

Which of the following is a typical goal of MAC spoofing?

Answers:

- ***Bypass 802.1x port-based security**
- Cause a switch to enter fail open mode
- Cause incoming packets to broadcast to all ports
- Reroute local switch traffic to a specified destination

Explanation:

MAC spoofing is changing the source MAC address on frames sent by the attacker. It is typically used to bypass 802.1x port-based security, to bypass wireless MAC filtering or hide the identity of the attacker's computer.

MAC flooding causes a switch to enter fail open mode, which causes incoming packets to be broadcast out to all ports. ARP spoofing/poisoning associates the attacker's MAC address with the IP address of the victim.

q_swi_attack_port_sec_secp8

Which protocol should you disable on the user access ports of a switch?

Answers:

- TCP
- ***DTP**
- PPTP
- IPsec

Explanation:

Switches have the ability to automatically detect ports that are trunk ports and to negotiate the trunking protocol used between devices. DTP is not secure and allows unauthorized devices to possibly modify configuration information. You should disable DTP services on the switch's end user (access) ports.

Transmission Control Protocol (TCP) is a communications standard that enables application programs and computing devices to exchange messages over a network and is important for traffic passing through a switch.

Point-to-Point Tunneling Protocol (PPTP) is a network protocol that enables the secure transfer of data from a remote client to a private enterprise server by creating a virtual private network (VPN) across TCP/IP-based data networks. PPTP can be important for the security of traffic passing through a switch.

IPSec is a set of communication rules or protocols for setting up secure connections over a network. Internet Protocol (IP) is the common standard that determines how data travels over the internet. IPSec adds encryption and authentication to make the protocol more secure, and can be important to for the security of traffic passing through a switch.

q_swi_attack_udp_secp8

You are a network security analyst for a large corporation. The company is planning to launch a new real-time video streaming service for its employees.

You've been tasked with selecting the most suitable transport protocol to ensure the best performance for this service.

Which of the following protocols would be the MOST effective for this purpose?

Answers:

- Transmission Control Protocol (TCP)
- ***User Datagram Protocol (UDP)**
- Secure Shell (SSH)
- Hypertext Transfer Protocol Secure (HTTPS)

Explanation:

User Datagram Protocol (UDP) is the correct answer. UDP is a connectionless protocol, making it faster than TCP and more suitable for real-time applications like video streaming, telephony, and gaming, where occasional packet loss is less impactful than latency. Therefore, UDP is the most effective choice for a real-time video streaming service.

Transmission Control Protocol (TCP) is a connection-oriented protocol that provides reliability, ordering, and error-checking. While it is suitable for applications requiring high levels of reliability, it is not the best choice for real-time applications like video streaming where occasional packet loss is less impactful than latency.

Secure Shell (SSH) is a cryptographic network protocol for operating network services securely over an unsecured network. While it provides secure data communication, it is not specifically designed for real-time applications like video streaming.

Hypertext Transfer Protocol Secure (HTTPS) is a secure version of HTTP, used primarily for secure communication over a computer network. While it provides secure communication, it is not specifically designed for real-time applications like video streaming.

5.10 Router Security

As you study this section, answer the following questions:

- Why should you change the default settings on new routers?
- Which secure protocols should you use to remotely manage a router?
- What actions can you take to ensure the physical security of network devices?
- Why should you update router firmware?
- How do ACLs work on a router?

In this section, you will learn to:

- Configure ACLs.
- Restrict Telnet and SSH access.
- Permit traffic.
- Block source hosts.

The key terms for this section include:

Term	Definition
Router	A network device that transmits data from one network to another.
Access control list (ACL)	A list of permissions associated with a network object, such as a router or a switch, that controls traffic at a network interface level.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.3 Explain various types of vulnerabilities.
	• Hardware
	  ○ Firmware
	
	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Access control
	  ○ Access control list (ACL)
	  ○ Permissions
	• Hardening techniques
	  ○ Disabling ports/protocols
	  ○ Default password changes
	
	3.2 Given a scenario, apply security principles to secure enterprise infrastructure.
	• Secure communication/access
	  ○ Remote access
	
	4.1 Given a scenario, apply common security techniques to computing resources.

	• Hardening targets
	  ○ Routers
	
	4.5 Given a scenario, modify enterprise capabilities to enhance security.
	
	• Implementation of secure protocols
	  ○ Protocol selection
	
	5.1 Summarize elements of effective security governance.
	
	• Standards
	  ○ Access control

TestOut Security Pro	2.2 Harden Network Devices 2.2.5 Configure Router Security

5.10.1 Router Security (Lesson Video)

Transcript:

We're going to talk about using a secure router configuration. If you're responsible for managing routers, it is critical that you use a secure router configuration. The way you do this depends on router manufacturer, model, configuration interface, and other factors.

In this lesson, we're going to look at some general router security principles that apply to a wide variety of router makes and models.
The most important aspect of securing any router, regardless of manufacturer, is changing default configuration settings. By default, routers are configured to use a predetermined username and password. Typically, this is something like admin and admin, or admin and password, or something very similar. There's actually even lists out on the internet that have the default username and password for practically every router model.
Because of this, it's extremely important you immediately change the default username and password on a router. In fact, I would recommend you changing this configuration before you even connect the router to the rest of the network or internet. This will ensure that there is zero chance someone will compromise the router before you have a chance to change this configuration.
When you're changing the default username, make sure to name it something other than admin or administrator or root or something easily guess. Also make sure you use a very complex password that includes numbers, symbols, upper- and lower-case letters, and is at minimum 12 characters long.
Another default setting you might consider changing is the default network address. Depending on how your organization is configured, you might have the ability to select a different IP address range than the default 192.168.1 network address that most routers use. This approach is known as security through obscurity, which by itself isn't actually secure, but if layered with other security techniques, it can add just one more obstacle an attacker has to get through, albeit a small one.
Depending on the device or organization you work for, you might be using a SOHO all-in-one router that includes switching and Wi-Fi capabilities. If this is the case, make sure to change the default SSID name and wireless encryption settings.
Another thing you should do when you first set up a router is make sure the firmware is up-to-date. Oftentimes routers are shipped with older firmware versions. Firmware updates contain much needed security fixes and vulnerability fixes, so it's important you are using the latest firmware version available. And always make sure you keep the router updated with the latest firmware.
The next configuration changes you want to make is that of the protocol used to connect to the router.

Some routers use connection protocols that send authentication credentials in cleartext. This makes it very easy for someone to capture packets and identify the login information for the device. While most modern routers will use some type of secure protocol, such as SSL, it's possible you might work with an older router that uses an older, unencrypted protocol by default. If this is the case, be aware that you should use a secure protocol.
For example, some routers might have FTP functionality that allows you to upload configuration files or firmware updates to the router. Because FTP is an unencrypted transmission protocol, the configuration file could be easily captured. Instead, you might be able to use SFTP, which is the secure, encrypted form of FTP.
Now, this does all depend on how you are connecting to the router to configure it. Some routers have a dedicated serial port that is used to configure the router. This is done instead of a remote configuration portal using a web browser or SSH. If this is the case, you might not need to worry about encrypted transmissions. However, if you are working with the router remotely, be sure you know which transmission protocol is being used and whether or not it is secure.
Speaking of using secure protocols, there's one protocol that I want to quickly talk about that some routers use, especially SOHO routers. That protocol is universal plug and play (UPnP).
UPnP was designed to make setting up media and IoT devices, such as smart TVs and gaming consoles, a breeze. It works by opening up ports for devices inside the network on the fly, eliminating the need to manually open up ports for these devices. In theory, it's a nice idea. However, in practice? Not so much.
UPnP has been plagued with exploit after exploit since it was implemented. These exploits range from annoying to downright scary. Because of this, I highly recommend simply disabling UPnP on any router that has it. At this time, the convenience isn't worth the risk.
Another area you need to look at with routers is that of remote access.
Most routers will allow you to configure remote access optionsâ€”both remote access from within the network and remote access from outside the network via the internet.
It might not be possible to disable remote access from inside the network. For example, the router might be locked inside of a networking closet or server rack case and accessing the router would require you to be inside the frigid server room. If this is the case, make sure the login credentials for remote access are configured properly and the protocol used is secure.
Remote access from outside, on the other hand, should 99% of the time be disabled. There are really only a handful of times remote access from outside the network would be necessary. And for these occasions, I recommend only enabling for as long as you need it, and not a moment longer. Allowing remote access from outside the network opens up entire attack surface that you shouldn't have to worry about.
Like I said earlier, you might not be able to have physical access to the router all the time, and this is actually a good thing.
As you already know, it's always a good idea to keep networking devices and critical infrastructure components under lock and key, and routers are no different. With physical access, it's possible for an attacker to access a router even if they don't have the proper login credentials. Because of this, make sure routers are stored in a secure location. At the very least, keep the router in a locked cabinet. But keeping it with other networking equipment in a locked cage or lock server room that has the appropriate physical security controls is best.
Another important aspect of proper router security is that of the configuration file.
The configuration file stores all the configure settings for the router, including open ports, usernames, firewall settings, et cetera. And often times these files can be backed up remotely on a USB drive or workstation. If you do backup these files, be sure to properly encrypt them so they can't easily read.
In addition, a lot of routers offer the option to encrypt the configuration file on the device itself. If your router happens to have this option, enable it. This provides an additional layer of security should someone gain access to this file.
And speaking of backing up the configuration file, do so on a regular basis. And as I said earlier, secure you backups by encrypting them.
Your network is only as strong as its weakest link. And if that weak link is the primary entry point into your network, as the router is, you might have a lot of problems on your hands. Keep this in mind when securing your network and take the extra time and steps we discussed in this lesson to properly secure your router.

5.10.2 Router ACLs (Lesson Video)

Transcript:

Like a firewall, a router can use access control lists (ACLs) to protect a network from attacks and control what type of communications are allowed on the network.

Don't confuse the ACL on a router with its routing table. The routing table is used to determine where network packets are forwarded to. The ACL, on the other hand, is used to determine whether or not a packet is allowed to be forwarded. For example, an ACL on a router could ban all incoming traffic destined to a specific IP address inside the network. Any packet destined for that IP address is dropped. However, outgoing packets from that IP address are still permitted.
When you create an ACL on a router, it almost always includes a hidden deny any statement at the end of the list. This means all traffic is automatically blocked. To allow traffic, you need at least one permit statement that either permits a specific traffic type or permits all traffic not specifically restricted.
The type of traffic you are able to control with an ACL depends on the type of ACL you create. Routers primarily use two ACL types: standard ACLs and extended ACLs.
A standard ACL is only able to filter traffic based on the source host name or host IP address--that is, it can't filter traffic based on port number, destination host name, or host IP address.
For example, you could configure a standard ACL to only allow traffic coming from internal IP addresses. All traffic coming from external IP address would be dropped. You could also configure a standard ACL to block all traffic from a specific IP address, but allow all other traffic. Know that because of how they function, standard ACLs should be placed as close to the destination as possible.
So, what if you wanted to block traffic based on other parameters, such as port number? To do this, you use the second type of ACLâ€""an extended ACL.
Extended ACLs are used to filter traffic based on a lot more parameters than standard ACLs. In addition to filtering based on source host name or host IP address, an extended ACL can filter based on source IP protocol, source or destination socket number, and destination host name or host IP address.
For example, you could block all outgoing traffic except for traffic on ports 80 and 443, or you could block all traffic from inside the network that's destined for a specific IP address.
Unlike standard ACLs, extended ACLs should be placed as close to the source as possible.
That's it for this lesson. In this lesson, we looked at router ACLs. We looked at how ACLs on a router work and what parameters are defined when they are created. We then looked at the two types of ACLs routers use, standard ACLs and extended ACLs.

5.10.3 Router Security Facts

5.10.4 Configuring ACLs (Demo Video)

Transcript:

In this demonstration, we'll show you how to configure access control lists on a Cisco router using the command line interface. Access control lists can give us granular control over how data flows across our networks, and which computers can communicate with specific protocols.

An access control list is simply a sequential list of statements of what can and cannot be done on the network. Those that can be done use the permit keyword. Those that can't be done use the deny keyword.
The ordering of the lines is important. These lists should go from the most specific conditions to the most general.
For instance, if we wanted to disallow FTP on our network, we could create a statement added to our access list that denied FTP. We would do that first, and then later we might allow all other application protocols that sit on top of the IP protocol.
And so, we might have two rules there. It's also important to recognize that there's an implicit deny anything else or any other traffic. This is located at the end of every access control list. In other words, anything that isn't explicitly permitted is implicitly denied.
For our demonstration, we're going to use this Windows 10 virtual machine that you can see over here. You can see that it currently has access to the Internet. We could look up access control lists and find out some information.
We can also see that we can currently ping, say, Google's DNS servers. We have access to the Internet. Let's go ahead and prohibit, or deny access for this computer to the Internet.
I'm going to attach to my Cisco router using Tera Term here. I have a cable going to the console port on the router. You can see that my router is currently, already configured to allow NAT, and it also has some specific configuration settings allowing different networks to connect with each other.
If we wanted to deny access to the Internet for this device, we need to know what its IP address is. We can get that by typing IPconfig here. You can see this is 10.0.0.2.

You can see that it must be connected to this 10.0.0.1 network that's associated with interface FA0/1. To deny a specific host, it's fairly straightforward. We need to get into configuration mode.
Now, you might have read, or you might know that there are two different types of access control lists, with Cisco. There are standard lists, and there are extended ACL. Standard lists only act on the source IP address for filtering.
You can only deny or permit specific hosts based on IP address, not protocols. Whereas, extended access control lists allow you to get pretty granular with which ports are allowed and denied, which protocols are allowed or denied also.
In our case, standard control list will work fine for this action. Now that we're in configuration mode, we're going to create an access list, and we'll just call it access list 10. Standard lists range typically from 1 to 99, or from 1300 to 1999. We'll just choose 10 because that says this is a standard access list.
And then, we will state that we want to deny 10.0.0.2, and specifically just that host. This second parameter is a wild card mask. Essentially, anything that's a zero there, says it has to match, or the IP address has to match that.
In this case, we're only applying this rule to 10.0.0.2. Everything has to match that.
We'll do that. Now, we still have access to the Internet. We can go ahead and refresh. You can see that we still get results. We can't just create the access list. We now need to apply that access list.
And so, in order to apply it, we basically tell it which interface this applies to. We're already in configuration mode, so we'll go to FA0/1, and we'll apply this access list to that.
We're going to apply access list 10 to this interface, and we're going to examine packets as they arrive at this interface on the router. We'll go ahead and say, "Enter there."
And now, if we come over here, we should see that we no longer have access to the Internet. It just kind of spins there. If we come up, and we try to ping, it's unreachable.
Now, one problem with this list is we explicitly denied this host. We've achieved what we wanted to. However, because there's that implicit deny at the end, we've denied access to everything else as well, everything that's coming through FA0/1.
If we wanted to keep that there, we'd want to, maybe create a rule that said access list 10 permit, and then the range of IP addresses that would be permitted.
In this case, let's just go ahead and just disable that access control list. We're still on that, so I'm going to simply type, no IP access group 10 in, which basically reverses the command that I had just done.
We should see that now we have access back out to the Internet. We can refresh this one, as well.
What if we wanted to deny ping, which uses ICMP, while still allowing Internet access? How can we get a little bit more granular? Well, we're going to need extended control lists in order to do that.
Let's go ahead and create an extended access control list. We'll get out of this interface. Let's go ahead and create access list 100. For this one, let's first deny ICMP, and we can specify a single host by just saying, "Host."
We'll say, "Deny this specific host for ICMP." And, to what network? We'll say, "To anything that comes through." The 0.0.0 says, "Anything." Coupled with this wildcard mask then that denies all ICMP.
Now, we haven't applied this list yet, so it still will allow us out to the Internet. We also want to permit access to other types of protocols, or any other protocol besides ICMP. Our second rule on our list will allow us to do that.
We're going to permit IP for that host. We can shorten it and just say, "To any," if we wanted to. Or, maybe we could be a little broader and say, "We're actually going to allow anyone on the 10 dot network to have access to any network."
Let's go ahead and try to apply that to FA0/1. Let's go ahead and try it. So, we are not able to ping, but we are able to get out to the Internet. Let's go ahead and show you that, and that it's not just cached. You can do different queries.
Access control lists are very powerful but can become complicated. It's important that you understand the basics that essentially these are just a sequential list of permit-and-deny statements, with an implicit deny at the end.
Standard access control lists allow you to filter based on a source IP address. Extended access control lists allow you to filter more granularly on specific protocols, or even specific ports for individual hosts, or for whole ranges of addresses.
Okay. That's it for this demonstration. In this demonstration, we showed you how to configure access control lists on a Cisco router using the command line interface.
Access control lists are foundational to allow us to screen the data that comes through our networks that we either allow out or into our corporate networks.

5.10.5 Restrict Telnet and SSH Access (Simulation)

Scenario

You are in the process of configuring a new router. The router interfaces connect to the following networks:

Interface	Network
FastEthernet0/0	192.168.1.0/24
FastEthernet0/1	192.168.2.0/24
FastEthernet0/1/0	192.168.3.0/24

Only Telnet and SSH access from these three networks should be allowed.

In this lab, your task is to:

- Use the **access-list** command to create a standard numbered access list using number 5.
- Add a **permit** statement for each network to the access list.
- Use the **access-class** command to apply the access list to VTY lines 0–4. Use the **in** direction to filter incoming traffic.
- Save your changes in the **startup-config** file.

Explanation

Complete this lab as follows:

1. Enter the configuration mode for the router:
    a. From the exhibit, select the router.
    b. From the terminal, press **Enter** .
    c. Type **enable** and then press **Enter** .
    d. Type **config term** and then press **Enter** .
2. From the terminal, create a standard numbered access list using number 5. Add a **permit** statement for each network to the access list.
    a. Type **access-list 5 permit 192.168.1.0 0.0.0.255** and then press **Enter** .
    b. Type **access-list 5 permit 192.168.2.0 0.0.0.255** and then press **Enter** .
    c. Type **access-list 5 permit 192.168.3.0 0.0.0.25** 5 and then press **Enter** .
3. Apply the access list to VTY lines 0–4. Filter incoming traffic.
    a. Type **line vty 0 4** and then press **Enter** .
    b. Type **access-class 5 in** and then press **Enter** .
    c. Press **Ctrl** + **Z** .
4. Save your changes in the **startup-config** file.
    a. Type **copy run start** and then press **Enter** .
    b. Press **Enter** to begin building the configuration.
    c. Press **Enter** .

5.10.6 Permit Traffic (Simulation)

Scenario

The Fiji router has been configured with Standard IP Access List 11. The access list is applied to the *Fa0/0 interface* . The access list must allow all traffic except traffic coming from hosts 192.168.1.10 and 192.168.1.12. However, you've noticed that it's preventing all traffic from being sent on Fa0/0. You remember that access lists contain an implied **deny any** statement. This means that any traffic not permitted by the list is denied. For this reason, access lists should contain at least one permit statement, or all traffic is blocked.

In this lab, your task is to:

- Add a **permit any** statement to Access List 11 to allow all traffic other than the restricted traffic.
- Save your changes in the **startup-config** file.

Explanation

Complete this lab as follows:

1. Enter the configuration mode for the Fiji router:
    a. From the exhibit, select the *Fiji router* .
    b. From the terminal, press **Enter** .
    c. Type **enable** and then press **Enter** .
    d. Type **config term** and then press **Enter** .
2. From the terminal, add a **permit any** statement to Access List 11 to allow all traffic other than the restricted traffic.
    a. Type **access-list 11 permit any** and press **Enter** .
    b. Press **Ctrl** + **Z** .
3. Save your changes in the **startup-config** file.
    a. Type **copy run start** and then press **Enter** .
    b. Press **Enter** to begin building the configuration.
    c. Press **Enter** .

5.10.7 Block Source Hosts (Simulation)

Scenario

You have a small business network connected to the internet through a single router as shown in the network diagram. You have noticed that three hosts on the internet have been flooding your router with unwanted traffic. As a temporary measure, you want to prevent all communication from these three hosts until the issue is resolved.

In this lab, your task is to:

- Create a Standard Access List 25.
- Add statements to the access list to block traffic from the following hosts:
    - **199.68.111.199**
    - **202.177.9.1**
    - **211.55.67.11**
- Add a statement to allow all other traffic from all other hosts.
- Apply Access List 25 to the *Serial0/0/0 interface* to filter incoming traffic.

You can also use **199.68.111.199 0.0.0.0** (without the **host** parameter) to identify a specific host. You can also use **0.0.0.0 255.255.255.255** to identify any host. Because this is a temporary solution, you do not need to save your changes.

Explanation

Complete this lab as follows:

1. Enter the configuration mode for the router:
    a. From the exhibit, select the router.
    b. From the terminal, press **Enter** .
    c. Type **enable** and then press **Enter** .

d. Type **config term** and then press **Enter** .

2. From the terminal, create a standard numbered access list using number 25. Add statements to the access list to block traffic to the required hosts.
    a. Type **access-list 25 deny host 199.68.111.199** and press **Enter** .
    b. Type **access-list 25 deny host 202.177.9.1** and press **Enter** .
    c. Type **access-list 25 deny host 211.55.67.11** and press **Enter** .
3. From the terminal, add a statement to allow all other traffic from all other hosts, by typing **access-list 25 permit any** and pressing **Enter** .
4. From the terminal, apply Access List 25 to the *Serial0/0/0 interface* to filter incoming traffic.
    a. Type **int s0/0/0** and press **Enter** .
    b. Type **ip access-group 25 in** and press **Enter** .
    c. Type **Ctrl** + **Z** .

5.10.8 Practice Questions (Section Quiz)

q_router_sec_acl_secp8

Which of the following should be configured on the router to filter traffic at the router level?

Answers:

- ***Access control list**
- Anti-spoofing rules
- SSH
- Telnet

Explanation:

Router access control lists (ACLs) can be configured to increase security and limit traffic, much like a firewall but on the router level. ACLs filter the traffic and determine if the data should be blocked or forwarded.

Anti-spoofing rules counter spoofing attacks where IP packets have a source address that does not belong to the sender.

Secure Shell (SSH) is a secure protocol that can be used to connect to the router.

Telnet is an older protocol used to connect to remote devices. It should not be used any longer.

q_router_sec_anti_spoofing_secp8

You are a network security engineer at a large corporation. You have been tasked with implementing anti-spoofing rules on the company's routers to enhance network security.

You have noticed an increase in spoofing attacks where IP packets have a source address that does not belong to the sender.

Which of the following anti-spoofing rule configurations would be the MOST effective in mitigating these attacks?

Answers:

- Source: Any, Destination: Any, Service: Any, Interface: Any external interface, Direction: Inbound, Action: Deny, Time: Any

- ***Source: An IP address belonging to the internal network or the IP address of the router itself, Destination: Any, Service: Any, Interface: Any external interface, Direction: Inbound, Action: Deny, Time: Any**
- Source: An IP address belonging to the external network, Destination: Any, Service: Any, Interface: Any internal interface, Direction: Outbound, Action: Deny, Time: Any
- Source: Any, Destination: An IP address belonging to the internal network or the IP address of the router itself, Service: Any, Interface: Any external interface, Direction: Outbound, Action: Deny, Time: Any

Explanation:

The correct answer is: Source: An IP address belonging to the internal network or the IP address of the router itself, Destination: Any, Service: Any, Interface: Any external interface, Direction: Inbound, Action: Deny, Time: Any. This configuration would effectively block any inbound packets that are pretending to be from the internal network or the router itself, which is a common tactic in spoofing attacks.

Source: Any, Destination: Any, Service: Any, Interface: Any external interface, Direction: Inbound, Action: Deny, Time: Any would block all inbound traffic, regardless of the source, which is not a practical solution as it would disrupt normal network operations.

Source: An IP address belonging to the external network, Destination: Any, Service: Any, Interface: Any internal interface, Direction: Outbound, Action: Deny, Time: Any would block all outbound traffic from any external IP address, which is not a practical solution as it would disrupt normal network operations and does not directly address the issue of spoofing attacks.

Source: Any, Destination: An IP address belonging to the internal network or the IP address of the router itself, Service: Any, Interface: Any external interface, Direction: Outbound, Action: Deny, Time: Any would block all outbound traffic destined for the internal network or the router itself, regardless of the source. While this might prevent some types of attacks, it does not specifically target spoofing attacks where the source address is falsified.

q_router_sec_blocked_secp8

Which of the following happens by default when you create and apply a new ACL on a router?

Answers:

- All traffic is permitted.
- ***All traffic is blocked.**
- The ACL is ignored until applied.
- ACLs are not created on a router.

Explanation:

When first created and applied on a router, an ACL almost always includes a hidden Deny Any statement at the end of the list. This means all traffic is automatically blocked.

All traffic is not permitted by default with a new ACL.

The ACL is immediately applied and blocks all traffic until configured.

ACLs are created on the router itself.

q_router_sec_firmware_secp8

You are deploying a brand new router. After you change the factory default settings, what should you do next?

Answers:

- Secure the configuration file.
- ***Update the firmware.**
- Configure SSH to access the router configuration.
- Configure anti-spoofing rules.

Explanation:

After changing the default settings on the router, you should update the firmware. Updates to the firmware fix any vulnerabilities that have been resolved by the manufacturer in the past.

After updating the firmware, you should configure the protocol used to connect to the router.

The configuration file stores all the configuration settings for the router, including open ports, usernames, firewall settings, and more. If possible, store the router configuration file in an encrypted form and back up the file to a secure location.

Anti-spoofing rules counter spoofing attacks where IP packets have a source address that does not belong to the sender. This is configured after the router is set up.

q_router_sec_locks_secp8

Which of the following can make passwords useless on a router?

Answers:

- ***Not controlling physical access to the router**
- Using the MD5 hashing algorithm to encrypt the password
- Storing the router configuration file in a secure location
- Using SSH to remotely connect to a router

Explanation:

If someone can gain access to the physical device, they can easily bypass any configured passwords. Passwords are useless if physical access is not controlled.

Other security measures you can use include:

- Use the MD5 hashing algorithm to encrypt the password.
- Store the router configuration file in encrypted form to a secure location.
- Use SSH when you remotely connect to a router.

q_router_sec_segmentation_01_secp8

A financial institution is processing transactions and wishes to improve its security posture. The institution divides its network into different sections to minimize risk while actively updating or retrieving transaction data.

What method does the financial institution intend to use?

Answers:

- ***Segmentation**

- Firewalling
- Network address translation (NAT)
- Virtual private network (VPN)

Explanation:

Segmentation works by dividing a network into different sections so that the institution can protect its data while processing.

Firewalling provides a line of defense against attacks but does not involve dividing the network into different sections.

NAT is a method of remapping one Internet Protocol (IP) address space into another by modifying network address information in the IP header of packets. However, it does not provide a method of dividing a network into segments.

A VPN can provide secure remote access to an organization's network, although it does not involve segmenting the network.

q_router_sec_segmentation_02_secp8

A critical infrastructure organization responsible for managing energy distribution across a large region relies heavily on industrial control systems (ICS) and supervisory control and data acquisition (SCADA) systems to monitor and control the power grid.

Given the critical nature of the operations, the IT team has implemented a unique control to safeguard these systems.

Which unique control did the IT team use to protect ICS and SCADA systems?

Answers:

- Regular system updates
- ***Network segmentation**
- Secure boot mechanisms
- Transport encryption protocols

Explanation:

Network segmentation isolates industrial control systems and supervisory control and data acquisition systems from the wider network, protecting them by limiting access and defending against inbound attacks.

Overall security benefits from regular system updates, but they are not unique controls designed to protect ICS/SCADA systems.

Secure boot mechanisms are focused on individual devices such as workstations and servers, not systems such as ICS and SCADA.

Transport encryption protocols are designed to secure the transmission of data, but are not as effective as network segmentation for isolating and protecting systems such as ICS and SCADA systems.

q_router_sec_segmentation_03_secp8

Upon reviewing the results of an organizational assessment, the cyber team implements various remediation practices to safeguard the company's data.

What remediation practices include the division of a network into separate pieces to contain an attack or attempted breach within one piece of the network?

Answers:

- ***Segmentation**
- Compensating controls
- Patching
- Insurance

Explanation:

In this instance, segmentation involves dividing a network into separate pieces to contain potential security breaches.

Not relevant here, compensating controls refer to measures put in place to mitigate the risk of a vulnerability when security teams cannot directly eliminate it or when direct remediation is not immediately possible.

Additionally, patching is one of the most straightforward and effective remediation practices. It involves applying updates and patches to software or systems to fix known vulnerabilities.

Insurance provides financial protection in case of a security breach resulting from a vulnerability and is important in a comprehensive risk management strategy, complementing technical controls with financial risk transfer.

q_router_sec_ssh_01_secp8

You've just deployed a new Cisco router that connects several network segments in your organization.

The router is physically located in a server room that requires an ID for access. You've backed up the router configuration to a remote location in an encrypted file. You access the router configuration interface from your notebook computer using a Telnet client with a username of admin and a password of P@ssW0rd. You have used the MD5 hashing algorithm to protect the password.

What should you do to increase the security of this device? (Select two.)

Answers:

- ***Use an SSH client to access the router configuration.**
- ***Change the default administrative username and password.**
- Use encrypted Type 7 passwords.
- Use a web browser to access the router configuration using an HTTP connection.
- Use TFTP to back up the router configuration to a remote location.

Explanation:

In this scenario, two key security issues need to be addressed. They are:

- You should use an SSH client to access the router configuration. Telnet transfers data in cleartext over the network connection, exposing sensitive data to sniffing.
- You should change the default administrative username and password. Default usernames and passwords are readily available from websites on the internet.

Encrypted Type 7 passwords on a Cisco device are less secure than those protected with MD5. Using HTTP and TFTP to manage the router configuration could expose sensitive information to sniffers, as these protocols transmit data in cleartext.

q_router_sec_ssh_02_secp8

You've just deployed a new Cisco router that connects several network segments in your organization.

The router is physically located in a locked server closet. You use an FTP client to regularly back up the router configuration to a remote server in an encrypted file. You access the router configuration interface from a notebook computer that is connected to the router's console port. You've configured the device with the username admin01 and the password P@ssW0rd. You have used the MD5 hashing algorithm to protect the password.

What should you do to increase the security of this device?

Answers:

- Use an SSH client to access the router configuration.
- Use encrypted Type 7 passwords.
- Move the router to a secure data center.
- ***Use SCP to back up the router configuration to a remote location.**

Explanation:

In this scenario, the router configuration is being copied to a remote location using an unsecure protocol (File Transfer Protocol) that transfers data in cleartext. You should instead use the Secure Copy Protocol (SCP) to transfer the backup from the router to the remote storage location.

It is not necessary to use an SSH client when using the console port to configure the router.

It is also not necessary to move the device to a data center if it is currently located in a locked server closet.

Encrypted Type 7 passwords on a Cisco device are less secure than those protected with MD5.

6.0 Resiliency and Site Security

6.1 Physical Threats

As you study this section, answer the following questions:

- What types of physical controls can be implemented to protect the perimeter of a building?
- How does an access control vestibule work?
- How are physical access controls similar to technical system security?
- How can environmental design enhance physical security?
- What are four types of sensors and how do they work?

In this section, you will learn to:

- Implement physical security.

The key terms for this section include:

Term	Definition
Physical security	Physical security is the protection of corporate assets from threats such as unauthorized entry, theft or damage.
Access list	A list of personnel who are authorized to enter a secure facility.
Access control vestibule	A specialized entrance with two locking doors that create a security buffer zone between two areas.
Bollard	Bollards are short, sturdy posts used to prevent a vehicle from crashing into a secure area.
Smart card	Access cards that have encrypted access information. Smart cards can be contactless or require contact.
Proximity card	Proximity cards, also known as radio frequency identification (RFID) cards, are a subset of smart cards that use the 125 kHz frequency to communicate with proximity readers.
Biometric Locks	Biometric locks increase security by using fingerprints or iris scans. They reduce the threat from lost keys or cards.
Sensor	A component in an alarm system that identifies unauthorized entry via infrared-, ultrasonic-, microwave-, or pressure-based detection of thermal changes or movement.

Radio frequency ID (RFID)	A means of encoding information into passive tags which can be energized and read by radio waves from a reader device.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts.
	• Physical security
	  ○ Bollards/barricades
	  ○ Access control vestibule
	  ○ Fencing
	  ○ Video surveillance
	  ○ Security guard
	  ○ Access badge
	  ○ Lighting
	  ○ Sensors
	  ▪ Infrared
	  ▪ Pressure
	  ▪ Microwave
	  ▪ Ultrasound
	2.4 Given a scenario, analyze indicators of malicious activity.
	• Physical attacks
	  ○ Brute force
	  ○ Radio frequency identification (RFID) cloning
	  ○ Environmental
	3.4 Explain the importance of resilience and recovery in security architecture.
	• Power
	4.6 Given a scenario, implement and maintain identity and access management.
	• Multifactor authentication
	  ○ Biometrics

TestOut Security Pro	2.1 Harden Physical Access 2.1.1 Implement physical security

6.1.1 Physical Security (Lesson Video)

Transcript:

Physical security is an integral aspect of cybersecurity operations, forming the first line of defense against threats targeting an organization's critical assets.

In cybersecurity, the goal is to safeguard digital assets and, equally importantly, the physical components housing these assets, such as servers, data centers, and crucial infrastructure. To effectively understand physical security in cybersecurity, we'll explore some practical examples.
Physical security uses tangible measures to protect digital assets. Access control mechanisms, such as biometric scanners, smart cards, key fobs, and surveillance systems comprising video cameras, motion sensors, and alarms constitute vital components. Biometric scanners, like fingerprint readers, ensure that only authorized personnel gain entry. Surveillance systems function as vigilant sentinels, detecting and alerting against potential threats.
Furthermore, backup power sources, redundant cooling systems, and fire suppression mechanisms play a pivotal role in securing data centers that ensure business continuity in the face of physical disruptions.
The site layout, like a battlefield strategy, arranges the terrain for maximum security.
Fencing acts as the outermost perimeter protection. Effective security fencing is transparent, robust, and resistant to climbing attempts. It serves as a visible boundary.
However, it may impart an intimidating appearance, which may not be ideal for organizations seeking a more welcoming facade.
Security lighting transforms the facility into a well-lit sanctuary, especially during nighttime hours.
It fosters a sense of safety and serves as a deterrent to potential intruders.
The lighting design takes into account overall illumination, specific areas for surveillance purposes, and the avoidance of shadowy or glaring spots.
Barricades help obstruct unauthorized access. They guide individuals through defined entry and exit points, each armed with authentication mechanisms.
In high-risk areas susceptible to attacks, barricades like bollards and security posts are deployed to prevent vehicles from approaching buildings at high speed.
Surveillance mechanisms complement these barricades, identifying any unauthorized attempts at penetration.
Access control vestibules, resembling double-lock security, regulate entry to secure areas. These vestibules comprise two interlocking doors or gates, allowing one person to pass at a time. The first door opens once access is granted via an access control system, such as a card reader or biometric scanner. Only when the first door is securely shut will the second door open, ensuring one-person access. This prevents unauthorized entry or tailgating, an effective security measure commonly seen in data centers, government facilities, and financial institutions.
Access badges replace physical keys with plastic cards embedded with magnetic strips, RFID chips, or NFC technology. These badges grant access through swiping, tapping, or proximity to a reader.
A Physical Access Control System, or PACS, administers this access, controlling entry to specific locations within a building or site. It combines hardware like card readers with software and a centralized control network, logging access activity for security audits and investigations.
In physical security, armed or unarmed, human security guards provide on-ground protection. They verify identification, monitor checkpoints, and apply expertise to potential breaches. Their presence acts as a visual deterrent.
Video surveillance offers cost-effective vigilance, although with certain limitations. Cameras record movements and access but with a slower response. However, they complement security guards and, in some cases, can serve as the primary surveillance method.
Alarms in physical security serve as early warning systems, supplementing other security controls. They detect and notify security personnel and occupants of potential threats, discouraging unauthorized access or criminal activity.
Alarm systems often integrate with access controls, cameras, and motion sensors for a more comprehensive approach.
That's it for this lesson. In this lesson, we discussed how physical security serves as the first line of defense against potential breaches, ensuring that only authorized personnel can access and manipulate network infrastructure. We went through several examples of physical security options like site layout, fencing and lighting to barricades, and entry and exit points. We then looked at how access control vestibules work. Next, we looked at the use of access badges, security guards, and cameras. We finished the lesson going over alarm systems and sensors. By safeguarding the tangible assets that house digital data, physical security forms an essential foundation for comprehensive network security strategies.

6.1.2 Physical Security Facts

Physical security concepts are a critical component of cybersecurity. Physical security measures, such as access control, video surveillance, and environmental controls, help protect an organization's physical assets, including servers, datacenters,

and other critical infrastructure, from unauthorized access, theft, or damage. Effective physical security practices help prevent unauthorized physical access to sensitive data or systems and reduce the risk of insider threats. Organizations must integrate physical security practices into their cybersecurity strategies to provide a layered approach to security and protect against cyber threats that exploit physical vulnerabilities.

This lesson covers the following topics:

- Physical security controls
- Physical security through environmental design
- Gateways and locks
- Access badges
- Security guards and cameras
- Alarm systems and sensors
- Physical attacks

Physical Security Controls

Physical security is critical to cybersecurity operations because it provides the first line of defense against physical access to an organization's critical assets. Cybersecurity is about securing digital assets and protecting the physical components that house those assets, such as servers, data centers, and other critical infrastructure.

Practical examples of physical security measures in cybersecurity operations include access control mechanisms, such as biometric scanners, smart cards, key fobs, and surveillance systems, including video cameras, motion sensors, and alarms. Additionally, backup power, redundant cooling, and fire suppression systems are critical components of physical security in data centers.

Physical access controls depend on the same access control fundamentals as technical system security:

- Authentication—Creates access lists and identifies mechanisms to allow approved persons through the barriers.
- Authorization—Creates barriers around a resource to control access through defined entry and exit points.
- Accounting—Records when entry/exit points are used and detects security breaches.

Physical security is often implemented by incorporating zones. Each zone is separated by its own barrier(s). One or more security mechanisms control entry and exit points through the barriers. Progression through each zone should be increasingly restrictive.

Physical Security Through Environmental Design

Physical security through environmental design is an approach to security that uses the built environment to enhance security and prevent crime. This approach designs physical spaces, buildings, and landscapes to promote non-obvious security features. Physical security via environmental design can be used in various settings, such as residential neighborhoods, commercial districts, schools, and public spaces. By incorporating these design principles, organizations can enhance security, deter criminal activity, and promote a sense of safety and well-being among users. Additionally, physical security via environmental design can be a cost-effective way to improve security because it easily incorporates design elements into new or existing structures at a low cost.

The following table includes some effective environmental design options to enhance physical security:

Device Type	Explanation

Barricades and Entry/Exit Points	A barricade is something that prevents access. As with any security system, no barricade is completely effective; a wall may be climbed, or a lock may be picked. The purpose of barricades is to channel people through defined entry and exit points. Each entry point should have an authentication mechanism so that only authorized persons are allowed through. Effective surveillance mechanisms ensure the detection of attempts to penetrate a barricade by other means. Physical sites at risk of a terrorist attack will use barricades such as bollards and security posts to prevent vehicles from speeding toward a building.
Fencing	The exterior of a building may be protected by fencing. Security fencing needs to be transparent (so guards can see any attempt to penetrate it), robust (so that it is difficult to cut), and secure against climbing (which is generally achieved by making it tall and possibly by using razor wire). Fencing is generally effective, but the drawback is that it gives a building an intimidating appearance. Buildings that are used by companies to welcome customers or the public may use more discreet security methods.
Lighting	Security lighting is enormously important in the perception that a building is safe and secure at night. Well-designed lighting helps to make people feel safe, especially in public areas or enclosed spaces, such as parking garages. Security lighting also acts as a deterrent by making intrusion more difficult and surveillance (whether by camera or guard) easier. The lighting design needs to account for overall light levels, the lighting of particular surfaces or areas (allowing cameras to perform facial recognition, for instance), and avoid areas of shadow and glare.
Bollards	Bollards are generally short vertical posts made of steel, concrete, or other similarly durable materials and installed at intervals around a perimeter or entrance. Sometimes bollards are non-obvious and appear as sculptures or as building design elements. They can be fixed or retractable, and some models can be raised or lowered remotely. Bollards can serve several purposes, such as protecting pedestrians from vehicular traffic, preventing unauthorized vehicle access, and providing perimeter security for critical infrastructure and facilities. They are often used to secure government buildings, airports, stadiums, store entrances, and other public spaces. By preventing vehicles from entering restricted areas, bollards can help mitigate the risks of vehicular attacks and accidents.

There may be a few options to adjust the site layout in existing premises. When faced with cost constraints and the need to reuse existing infrastructure, incorporating the following principles can be helpful:

- Locate secure zones, such as equipment rooms, as deep within the building as possible, avoiding external walls, doors, and windows.
- Position public access areas so that guests do not pass near secure zones. Security mechanisms in public areas should be highly visible to increase deterrence.
- Use signage and warnings to enforce the idea that security is tightly controlled. Beyond basic "no trespassing" signs, some homes and offices also display signs from the security companies whose services they are using. These may convince intruders to stay away.
- Conversely, entry points to secure zones should be discreet. Do not allow an intruder the opportunity to inspect security mechanisms protecting such zones (or even to know where they are). Use industrial camouflage to make buildings and gateways protecting high-value assets unobtrusive or create high-visibility decoy areas to draw out potential threat actors.
- Try to minimize traffic passing between zones. The flow of people should be "in and out" rather than "across and between."

- Give high-traffic public areas high visibility to hinder the covert use of gateways, network access ports, and computer equipment, and simplify surveillance.
- In secure zones, position display screens or input devices away from pathways or windows. Use one-way glass only visible from the inside out so no one can look in through windows.

Gateways and Locks

To secure a gateway it must be fitted with a lock. A secure gateway will normally be self-closing and self-locking rather than dependent on the user to close and lock it. Lock types can be categorized as follows:

- Physical—are conventional locks that prevent the door handle from being operated without using a key. More expensive types offer greater resistance against lock picking.
- Electronic—are locks, rather than a key, that operate by entering a PIN on an electronic keypad. This type of lock is also referred to as cipher, combination, or keyless. A smart lock may be opened using a magnetic swipe card or feature a proximity reader to detect the presence of a physical token, such as a wireless key fob or smart-card.
- Biometric—is a lock integrated with a biometric scanner.

An access control vestibule is a security measure that regulates entry to a secure area. It involves two doors or gates that interlock and permit only one individual to pass through at a time. The first door opens after the person is granted access via an access control system, such as a card reader or biometric scanner. Once the person enters the vestibule, the first door shuts. The second door opens only when the first door is firmly shut. This guarantees only one person can enter or exit at a time, preventing unauthorized access or tailgating. Access control vestibules are frequently used in high-security settings like data centers, government buildings, and financial institutions to offer an additional layer of physical security control. They effectively deter unauthorized access to secure areas and safeguard sensitive assets against potential physical attacks.

Cable locks attach to a secure point on the device chassis. A server chassis might come with both a metal loop and a Kensington security slot. As well as securing the chassis to a rack or desk, the position of the secure point prevents the chassis from being opened without removing the cable first.

Access Badges

Access badges are a fundamental component of physical security in larger organizations where control over access to various locations is critical. Plastic cards embedded with magnetic strips, radio frequency identification (RFID) chips, or near - field communication (NFC) technology are issued to authorized individuals, such as employees, contractors, or visitors instead of physical keys. Access badges replace physical keys but provide access similarly. This is achieved by requiring the badge to be swiped, tapped, or brought into proximity with a reader at the access point, like a door or turnstile. The reader communicates with a control system to verify the badge's authenticity and the level of access granted to the badge holder. If the system recognizes the badge as valid and authorized for that area, the door unlocks, granting access.

It is important to note that implementing this type of access control system requires magnetic door-locking mechanisms and access card readers, which depend upon electrical power and network communications at each access point (such as a doorway.)

A physical access control system (PACS) is a critical component in managing and maintaining security within a facility. It is a system designed to control who can access specific locations within a building or site. The PACS operates through a combination of hardware and software, including access cards or badges, card readers, access control panels, and a centralized control network. The PACS system provides valuable badge access activity logging capabilities.

In addition to controlling access, access badges also serve as a form of identification, displaying pertinent information about the badge holder, such as their name, title, and photograph. This aids in quickly identifying individuals within a facility and verifying that they are in an area appropriate for their role or purpose. Moreover, access badges can provide valuable data for security audits and investigations. Each time a badge is used, a PACS system can log the time, location, and identity

associated with the access event. This can be crucial in investigating security breaches, understanding movement patterns, and even planning emergency evacuation strategies.

Security Guards and Cameras

Surveillance is typically a second layer of security designed to improve the resilience of perimeter gateways. Surveillance may be focused on perimeter areas or within security zones. Human security guards, armed or unarmed, can be placed in front of and around a location to protect it. They can monitor critical checkpoints and verify identification, allow or disallow access, and log physical entry events. They also provide a visual deterrent and can apply their knowledge and intuition to potential security breaches. The visible presence of guards is a very effective intrusion detection and deterrence mechanism, but it is correspondingly expensive. It may not be possible to place security guards within certain zones because they cannot be granted the appropriate security clearance. Training and screening of security guards is imperative.

Video surveillance is a cheaper means of providing surveillance than maintaining separate guards at each gateway or zone, though it is still not cheap to set up if the infrastructure is not already in place on the premises. It is also quite an effective deterrent. The other big advantage is that movement and access can be recorded. The main drawback compared to the presence of security guards is that response times are longer, and security may be compromised if not enough staff are in place to monitor the camera feeds.

The cameras in a CCTV network are typically connected to a multiplexer using coaxial cabling. The multiplexer can then display images from the cameras on one or more screens, allow the operator to control camera functions and record the images to tape or hard drive. Newer camera systems may be linked in an IP network using regular data cabling.

Camera systems and robotics can use AI and machine learning to implement smart physical security:

- Motion Recognition—Occurs when the camera system is configured with gait identification technology. This means the system can generate an alert when anyone within sight of the camera moves in a pattern that does not match a known and authorized individual.
- Object Detection—Occurs when the camera system can detect changes to the environment, such as a missing server or unknown device connected to a wall port.
- Drones/UAV—cameras mounted on drones can cover wider areas than ground-based patrols.

Alarm Systems and Sensors

Alarms play a vital role in physical security by supplementing other security controls. Alarms alert security personnel and building occupants of potential threats or breaches. They are both detective and deterrent controls, notifying of trouble and discouraging unauthorized access and criminal activity. Alarms are often integrated with other physical security controls, such as access control systems, surveillance cameras, or motion sensors, to enhance their effectiveness. The following list describes several common types of alarms:

- Circuit—Uses a circuit-based alarm that sounds when the circuit is opened or closed, depending on the type of alarm. For example, this could be a door or window opening or by a fence being cut. A closed-circuit alarm is more secure because it cannot be defeated by cutting the circuit like an open-circuit alarm.
- Motion Detection—Uses a motion-based alarm linked to a detector that is triggered by any movement within an area such as a room (defined by the sensitivity and range of the detector). The sensors in these detectors are either microwave radio reflection (similar to radar) or passive infrared (PIR), which detect moving heat sources.
- Noise Detection—Uses an alarm triggered by sounds picked up by a microphone. Modern AI-backed analysis and identification of specific types of sound can render this type of system less prone to false positives.
- Proximity—Uses radio frequency ID (RFID) tags and readers to track the movement of tagged objects within an area. This allows an alarm system to detect whether someone is trying to remove equipment.
- Duress—Uses an alarm triggered manually by staff if they come under threat. There are many ways of implementing this type of alarm, including wireless pendants, concealed sensors or triggers, and DECT

handsets or smartphones. Some electronic entry locks can also be programmed with a duress code different from the ordinary access code. This will open the gateway but also alert security personnel that the lock has been operated under duress.

Circuit-based alarms are suited for use at the perimeter and on windows and doors. These may register when a gateway is opened without using the lock mechanism properly or when a gateway is held open for longer than a defined period. Motion detectors are useful for controlling access to spaces not normally used. Duress alarms are useful for exposed staff in public areas. An alarm might simply sound like an alert or be linked to a monitoring system. Many alarms are linked directly to local law enforcement or third-party security companies. A silent alarm alerts security personnel rather than sounding an audible alarm.

Sensor Types

Sensors are critical in implementing physical security measures, providing proactive detection and alerting capabilities against potential security breaches. These devices can employ various technologies, including infrared, pressure, microwave, and ultrasonic systems, each with unique advantages and suitable applications.

- Infrared sensors are commonly used in motion detection systems. They detect changes in heat patterns caused by moving objects, such as a human intruder. These are often used in residential and commercial security systems, triggering alarms or activating security lights when detecting motion.
- Pressure sensors are typically installed inside floors or mats and are activated by weight. They can be used in high-security areas to detect unauthorized access or even in retail environments to count foot traffic.
- Microwave sensors emit microwave pulses and measure the reflection off a moving object. They are often combined with infrared detectors in dual-technology motion sensors. These sensors are less likely to trigger false alarms, as the infrared and microwave sensors must be tripped simultaneously to trigger an alarm. These can be useful in securing large outdoor areas like parking lots or fenced areas.
- Ultrasonic sensors emit sound waves at frequencies above the range of human hearing and measure the time it takes for the waves to return after hitting an object. They are often used in automated lighting systems to switch lights on when someone enters a room and switch them off again when the room is empty.

Physical Attacks

A physical attack is one directed against cabling infrastructure, hardware devices, or the environment of the site facilities hosting the network.

The following table include three of the most common physical attacks.

Attack	Explanation
Brute Force	A brute force physical attack can take several different forms, some examples of which are the following: • Smashing a hardware device to perform physical denial of service (DoS). • Breaking into premises or cabinets by forcing a lock or gateway. This is likely to be an indicator of theft or tampering. Preventing theft is often impossible to guarantee, so knowing that something has been stolen is important for things like data breach reporting and revoking access permissions. A system that is tamper-evident will display visible signs of forced entry or use that are difficult for a threat actor to disguise.

Environmental	An environmental attack could be an attempt to perform denial of service. For example, a threat actor could try to destroy power lines, cut through network cables, or disrupt cooling systems. Alternatively, environmental and building maintenance systems are known vectors for threat actors to try to gain access to company networks. The risk from physical attacks means that premises must be monitored for signs of physical damage or the addition of rogue devices.
RFID Cloning	Radio Frequency ID (RFID) is a means of encoding information into passive tags. When a reader is within range of the tag,it produces an electromagnetic wave that powers up the tag and allows the reader to collect information from it. This technology can be used to implement contactless building access control systems. RFID cloning and skimming refer to ways of counterfeiting contactless building access cards or badges: • Cloning —this refers to making one or more copies of an existing card. A lost or stolen card with no cryptographic protections can be physically duplicated. Card loss should be reported immediately so that it can be revoked and a new one issued. If there were a successful attack, it might be indicated by use of a card in a suspicious location or time of day. • Skimming —this refers to using a counterfeit reader to capture card or badge details, which are then used to program a duplicate. Some types of proximity card can quite easily be made to transmit the credential to a portable RFID reader that a threat actor could conceal on their person. These attacks can generally only target "dumb" access cards that transfer static tokens rather than perform cryptoprocessing. If use of the cards is logged, compromise might be indicated by impossible travel and concurrent use access patterns. Near-field communication (NFC) is derived from RFID and is also often used for contactless cards. It works only at very close range and allows two-way communications between NFC peers.

6.1.3 Implement Physical Security (Simulation)

Scenario

Based on a review of physical security at your office, you have recommended several improvements. Your plan includes installing smart card readers, IP cameras, signs, and an access log book.

In this lab, your task is to:

Implement your physical security plan by dragging the correct items from the shelf onto the various locations in the building. As you drag the items from the shelf, the possible drop locations are highlighted. To implement your plan, you must:

- Install two IP security cameras in the appropriate location to record which employees access the key infrastructure. The security cameras should operate over the TCP/IP network.
- Install the smart card key readers in the appropriate locations to control access to key infrastructure. The key card readers should be contactless and record more information than the card's ID.
- Install a Restricted Access sign on the networking closet door to control access to the infrastructure.
- Place the visitor log on the lobby desk.

Not all items on the shelf will be used.

Explanation

Complete this lab as follows:

1. Install the IP security cameras:
    a. From the Shelf, expand **CCTV Cameras** .
    b. Drag the IP Security Camera from the Shelf to the highlighted circle inside the networking closet.
    c. Drag the IP Security Camera from the Shelf to the highlighted circle just outside the networking closet.
2. Install the smart card key readers:
    a. From the Shelf, expand **Door Locks** .
    b. Drag a smart card reader from the Shelf to the highlighted location outside the building's front door.
    c. Drag a smart card reader from the Shelf to the highlighted location outside the networking closet's door.
3. Install the Restricted Access sign:
    a. From the Shelf, expand **Restricted Access Signs** .
    b. Drag the Restricted Access sign from the Shelf to the networking closet door.
4. Place the visitor log on the lobby desk:
    a. From the Shelf, expand **Visitor Logs** .
    b. Drag the visitor log from the Shelf to the lobby desk.

6.1.4 Practice Questions (Section Quiz)

q_phys_sec_badges_01_secp8

Which of the following are solutions that address physical security? (Select two.)

Answers:

- ***Require identification and name badges for all employees.**
- Implement complex passwords.
- Disable guest accounts on computers.
- ***Escort visitors at all times.**
- Scan all floppy disks before use.

Explanation:

Physical security controls physical access to the network or its components. Physical security controls include:

- Requiring identification or key cards before entry is permitted.
- Escorting visitors at all times.
- Keeping doors and windows locked.
- Keeping devices with sensitive information out of view of public users.
- Keeping the server room locked (locking computers to racks or tables to prevent theft).

Implementing complex password, disabling guest accounts on computers, and scanning all floppy disks before use are not considered physical security controls.

q_phys_sec_badges_02_secp8

You are the security administrator for a small business. The floor plan for your organization is shown in the figure below.

You've hired a third-party security consultant to review your organization's security measures. They have discovered multiple instances where unauthorized individuals have gained access to your facility, even to very sensitive areas. They recommend that you provide employees with access badges and implement access badge readers to prevent this from happening in the future.

Click on the office locations where access badge readers would be most appropriate.

Answers:

- ***162,137,131,132**
- ***315,40,112,104**
- 22,40,112,104
- 135,20,179,114
- 22,144,139,114
- 294,144,139,114

Explanation:

Access badge readers are typically implemented at building entrances to control access to a facility. Only individuals who have an authorized access badge are allowed to enter the facility. Individuals who do not have an access badge must be cleared and admitted by security personnel. Additional access badge readers can be implemented within the facility to further restrict access to sensitive areas, such as the server room.

q_phys_sec_biometric_secp8

If a fingerprint or retina scan is required to open a secured door, which kind of physical security has been implemented?

Answers:

- Access list
- Mantrap
- ***Biometric locks**
- Double-entry door

Explanation:

Biometric locks use unique physical characteristics of a person to authenticate his or her access to a secured item. Often, these locks take the form of fingerprint scanners or retina scanners.

An access list is incorrect because it is a list of names that a guard checks.

Mantraps and double-entry doors are also incorrect because they are styles of entryways and don't check physical characteristics.

q_phys_sec_bollards_secp8

A security manager decides to enhance the physical security of a warehouse storing high-value tech equipment by installing a deterrent at the perimeter to prevent vehicle-based attacks.

Which security measure would be the MOST suitable for this purpose?

Answers:

- ***Bollards**
- Access badge
- Access control vestibule
- Fencing

Explanation:

Bollards are short vertical posts that serve as an effective barrier against vehicle-based attacks. Their sturdy design stops vehicles from entering sensitive areas, making them the most suitable choice for preventing vehicle-based attacks.

An access badge does not serve to deter vehicle-based attacks. Its primary function is to verify the identity of individuals and grant or deny access based on their authorization levels.

An access control vestibule effectively controls pedestrian access but does not provide any substantial deterrent against vehicle-based attacks.

Fencing can serve as a physical barrier to prevent unauthorized access. However, compared to bollards, it is less effective against vehicle-based attacks, especially if the vehicles are large or moving at high speed.

q_phys_sec_cctv_01_secp8

Which option is a benefit of CCTV?

Answers:

- Reduce the need for locks and sensors on doors.
- ***Expand the area visible by security guards.**
- Provide a corrective control.
- Increase security protection throughout an environment.

Explanation:

A primary benefit of CCTV is that it expands the area visible by security guards. This helps few guards oversee and monitor a larger area.

CCTV does not reduce the need for locks and sensors on doors and does not provide a corrective control (it is a preventative, deterrent, or detective control). CCTV does not increase security protection throughout an environment, as the range is limited to areas over which it is aimed.

q_phys_sec_cctv_02_secp8

You want to use CCTV to increase your physical security, and you want the ability to remotely control the camera position.

Which camera type should you choose?

Answers:

- ***PTZ**

- Bullet
- C-mount
- Dome

Explanation:

A Pan Tilt Zoom (PTZ) camera lets you dynamically move the camera and zoom in on specific areas (cameras without PTZ capabilities are set looking a specific direction). Automatic PTZ mode automatically moves the camera between several preset locations. Manual PTZ lets an operator remotely control the position of the camera.

A bullet camera has a built-in lens and is long and round in shape. Most bullet cameras can be used indoors or outdoors. A c-mount camera has interchangeable lenses, is typically rectangular in shape, and carries the lens on its end. Most c-mount cameras require special housing to be used outdoors. A dome camera is a camera protected with a plastic or glass dome. These cameras are more vandal-resistant than other cameras.

Bullet, c-mount, or dome cameras can also be PTZ cameras.

q_phys_sec_cctv_03_secp8

A tech company recently moved to a new facility and seeks to bolster its physical security posture. The security team proposes integrating security guards and surveillance cameras as part of the security measures.

The chief security officer (CSO) wants to ensure these implementations effectively deter, detect, and report potential security incidents.

Given the scenario, which actions will maximize the effectiveness of security guards and cameras in enhancing the organization's physical security? (Select two.)

Answers:

- ***Position cameras to monitor critical access points and sensitive areas.**
- Allow security guards to monitor camera feeds only during break times.
- ***Implement security guard rotations and unannounced spot checks.**
- Place cameras in highly visible areas, but do not connect them to any recording device.
- Place two cameras on the locations where there are security guards positioned.

Explanation:

Positioning cameras to monitor critical access points and sensitive areas ensures comprehensive surveillance, enhancing deterrence and detection capabilities.

Implementing security guard rotations and unannounced spot checks prevents complacency and ensures guards remain vigilant, enhancing the detection and reporting of incidents.

Allowing security guards to monitor camera feeds only during break times reduces the effectiveness of real-time surveillance, making this measure counterproductive.

While visible cameras can act as a deterrent, not connecting them to any recording device negates their primary purpose of recording potential incidents for later review or evidence collection.

Placing cameras (one or more) on the locations where security guards are positioned, does not necessarily enhance physical security, as the security guards already provide significant security. It would be wiser to place two cameras in areas where there are no security guards located.

q_phys_sec_door_secp8

Which of the following controls is an example of a physical access control method?

Answers:

- ***Locks on doors**
- Access control lists with permissions
- Passwords
- Smart cards
- Hiring background checks

Explanation:

Locks on doors are an example of a physical access control method. Physical controls restrict or control physical access.

Passwords, access control lists, and smart cards are all examples of technical controls.

Even though a smart card is a physical object, the card by itself is a part of a technical implementation.

Requiring background checks for hiring is an example of a policy or an administrative control.

q_phys_sec_fencing_secp8

A data center must enhance its security measures to prevent unauthorized access to its facility. The center are considering different methods to achieve this goal.

What should the data center implement first to ensure a strong physical barrier against intrusions?

Answers:

- ***Fencing**
- Biometric authentication
- Security guard patrols
- Video surveillance

Explanation:

Fencing serves as a first line of defense in physical security, creating a physical barrier that prevents or discourages unauthorized access to the facility. It not only restricts intrusions but also acts as a psychological deterrent.

Biometric authentication effectively controls access inside the data center, although it is not a primary physical barrier against intrusions from outside.

Security guard patrols are not considered a physical barrier and thus do not provide the initial deterrence that fencing does.

Video surveillance is not a physical barrier to entry and therefore does not provide the initial deterrence that fencing offers.

q_phys_sec_gateways_locks_secp8

A major technology company plans to renovate its headquarters, emphasizing both physical and digital security.

The head of the security department is looking to enhance the building's main entry points and contemplates integrating advanced gateways with innovative locking mechanisms.

In relation to securing a major technology company's main entry points, which approaches will BEST harness the potential of gateways and locks to ensure optimal security? (Select two.)

Answers:

- ***Employ network gateways that scrutinize incoming traffic for malicious activity.**
- ***Implement biometric locks that grant access based on unique physiological characteristics.**
- Use gateways to redirect all visitors to the company's promotional website.
- Install traditional padlocks that require a standard key.
- Install CCTV cameras to increase the effectiveness of the gateways and locks.

Explanation:

Network gateways that evaluate incoming traffic for potentially harmful activity effectively act as an organization's first line of digital defense, thereby safeguarding the company's internal network resources.

Biometric locks enhance physical security by leveraging unique physiological traits, such as fingerprints or retina scans, thus reducing the likelihood of unauthorized access.

Redirecting visitors to a promotional website does not improve security. Instead, gateways should prioritize security over promotional activities.

Traditional padlocks with standard keys do not offer the advanced security necessary for a major technology company, as potential threat actors can easily duplicate or pick them.

Installing CCTV cameras can help in monitoring secured areas, but it does not necessarily increase the effectiveness of the gateways and locks.

q_phys_sec_infrared_sensor_secp8

To increase the physical security of a secured location, an organization deploys motion detection sensors throughout the grounds and building.

What type of sensor uses this technology?

Answers:

- ***Infrared sensor**
- Pressure sensor
- Microwave sensor
- Ultrasonic sensor

Explanation:

In this instance, the organization would install infrared sensors commonly used in motion detection systems. They can detect changes in heat patterns caused by moving objects, such as a human intruder.

While effective but not used to detect motion, pressure sensors activate when applying weight to the sensor.

Microwave sensors emit microwave pulses and measure the reflection of a moving object. However, they would also need infrared sensors to trigger simultaneously to use as motion detection.

Not ideal in this situation, ultrasonic sensors emit sound waves at frequencies above the range of human hearing and measure the time it takes for the waves to return after hitting an object.

q_phys_sec_layered_secp8

The cybersecurity team at a multinational corporation is collaborating with the facilities department to design a new data center. The team seeks to integrate top-tier physical security controls into the site layout to maximize protection against potential threats.

The discussions revolve around the BEST strategies to ensure the safety of the data center.

When designing the physical security controls for the site layout of the new data center, which strategy would be MOST effective in deterring unauthorized access and providing a comprehensive security layer?

Answers:

- ***Establishing a security perimeter with layered access controls**
- Implementing a single, fortified main entrance
- Placing all servers near windows for easy maintenance
- Distributing security personnel evenly throughout the premises

Explanation:

Layered access controls, often referred to as "defense in depth," ensure that multiple levels of security mechanisms are in place. If one fails or gets bypassed, others still provide protection.

While a fortified entrance can be a strong deterrent, relying solely on one point of entry without layered controls can lead to vulnerabilities. Single points of failure are generally not advisable.

Positioning servers near windows can expose them to potential threats, such as visual hacking or physical breaches. It is counterproductive from a security perspective.

While security personnel are important, their distribution should be strategic, focusing on critical points rather than a mere even spread. Layered access controls would still be more effective.

q_phys_sec_rfid_cloning_secp8

As the head of physical security at a large tech company, you have been tasked with investigating a series of unauthorized entries into secure areas of your facilities.

The intrusions have been sporadic and seemingly random, with no clear pattern or motive. The intruders have not been caught on camera, and no physical damage or theft has been reported. However, you notice that the access logs show entries made using the credentials of employees who were not on-site at the time of the incidents.

Which of the following is the MOST likely method the intruders are using to gain access?

Answers:

- Lock picking

- Social engineering
- ***RFID cloning**
- Bypassing CCTV cameras

Explanation:

RFID cloning is the most likely method used by the intruders. RFID cloning involves copying the data from an employee's RFID access card and using it to create a duplicate card. This would explain the access logs showing entries made using the credentials of employees who were not on-site at the time, as the intruders could be using cloned cards to gain access.

Lock picking is unlikely to be the method used as the access logs show entries made using employee credentials. Lock picking would not leave such a trace.

While social engineering could potentially explain the unauthorized entries, it would not account for the access logs showing entries made using the credentials of employees who were not on-site at the time.

While bypassing CCTV cameras could potentially allow the intruders to avoid detection, it would not explain the access logs showing entries made using the credentials of employees who were not on-site at the time.

q_phys_sec_smartcards_secp8

Which kind of access control technology allows more than just the identity of an individual to be transmitted wirelessly to either allow or deny access?

Answers:

- Proximity card
- ***Smart card**
- Biometric locks
- Keypad locks

Explanation:

Unlike proximity cards that only transmit the owner's identity, smart cards can contain and transmit many more pieces of information.

Biometric locks and keypad locks don't transmit data wirelessly. In contrast, they require physical interaction.

q_phys_sec_vestibule_secp8

A company wants to improve the physical security at its headquarters. They need a solution that can help regulate access to the building and deter potential intruders during nighttime.

Which physical security measure should they prioritize?

Answers:

- ***Access control vestibule**
- Enhanced lighting
- Closed-circuit television (CCTV)
- Perimeter fencing

Explanation:

An access control vestibule is a two-door system where the first door must close before the second opens, effectively controlling and managing access to the building.

Lighting can deter potential intruders during nighttime by illuminating dark areas, but it does not regulate access to the building.

CCTV can provide surveillance and record activity around the building, but it does not regulate access to the building. CCTV provides a way to monitor and record activities but does not control who can enter the building.

Perimeter fencing can deter potential intruders and control access to a property to an extent. However, it does not regulate access to the building in a controlled manner as an access control vestibule does.

6.2 Monitoring and Reconnaissance

As you study this section, answer the following questions:

- What is the goal of network monitoring?
- What is the difference between passive and active reconnaissance?
- What tool is a search engine for internet-connected devices?
- What are common techniques used in active reconnaissance?

In this section, you will learn to:

- Perform port and ping scans.
- Perform reconnaissance with Nmap.

The key terms for this section include:

Term	Definition
IP scanners	Special tools that allow a network administrator to scan the entire network to find all connected devices and their IP addresses.
Reconnaissance	Also known as *footprinting* . This is the process of gathering information about a target before beginning any penetration test or security audit.
Active reconnaissance	The process of gathering information by interacting with the target in some manner.
Passive reconnaissance	The process of gathering information about a target with no direct interaction with the target.
Packet sniffing	The act of capturing data packets transmitted across the network and analyzing them for important information.

War driving	The act of driving around with a wireless device looking for open vulnerable wireless networks.
War flying	The act of using drones or unmanned aerial vehicles to find open wireless networks.
Eavesdropping	The act of covertly listening in on a communication between other people.
Open-Source Intelligence (OSINT)	Any data that is collected from publicly available sources such as social media, search engines, company websites, media sources, or public government sources.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.2 Explain common threat vectors and attack surfaces.
	• Unsecure networks
	• Open service ports
	
	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Monitoring
	
	4.1 Given a scenario, apply common security techniques to computing resources.
	• Monitoring
	
	4.3 Explain various activities associated with vulnerability management.
	• Identification methods
	  ○ Open-source intelligence (OSINT)
	
	4.4 Explain security alerting and monitoring concepts and tools.
	• Monitoring computing resources
	  ○ Systems
	  ○ Infrastructure
	  ○ Infrastructure
	• Tools
	  ○ Vulnerability scanners
	
	4.5 Given a scenario, modify enterprise capabilities to enhance security.
	• Firewall
	  ○ Ports/protocols

	4.9 Given a scenario, use data sources to support an investigation.
	• Data sources
	  ○ Packet captures
	5.4 Summarize elements of effective security compliance.
	• Compliance monitoring
	5.5 Explain types and purposes of audits and assessments.
	• Internal
	  ○ Compliance
	• Penetration testing
	  ○ Reconnaissance
	  ▪ Passive
	  ▪ Active

TestOut Security Pro	2.2 Harden network devices 2.2.4 Harden a wireless network

6.2.1 Network Monitoring (Lesson Video)

Transcript:

Keeping an eye on the type of traffic that flows across your network helps you to better understand what's going on in your system. If you know your network and constantly monitor it, you can catch attackers before they do any significant damage. In this lesson, we'll look at some common network monitoring tools and how they can be used.

One of the more common network tests is a connection test. The command line utility named ping is the easiest way to do this. The ping utility sends an ICMP packet from our computer to another device on our network and waits for a response. If we get a response, the device is connected and active. If there's no response, it means the device could be offline or something is blocking the reply packets, and we need to figure out why. Ping is used in pretty much all operating systems.
Traceroute is a tool that allows us to see the path that a packet takes. Every device the packet travels through to its destination is known as a hop. The traceroute utility helps us determine if there's a specific device in our network that's down and causing problems for us. Traceroute is the command used in Linux and tracert is the command for Windows. The Linux version defaults to sending UDP packets and Windows uses ICMP packets. Other than that, these tools do exactly the same thing.
Pathping is a Windows command line tool that essentially combines the ping and tracert tools. The pathping utility sends ICMP packets to check connectivity, but it also gives us information such as latency and packet loss about each and every hop a packet takes. We can also use pathping to see if certain router connections are dropping packets before they reach their desired destination.
The netstat command shows us all sorts of network statistics. When we run this command, we receive a snapshot of our TCP connections, which ports they're connected to, the programs that are using those connections, and other statistics. We can use this information to monitor our traffic and make sure that our computer isn't connected to places it shouldn't be.
The route command is used in both Linux and Windows to show us our routing table and to allow us to make manual changes to it. Routers use the routing table to determine where to send packets. If there's an issue with the routing table, the route command can help us find it and make the necessary changes.

The arp command is another command used in basically all operating systems. ARP stands for Address Resolution Protocol, and its purpose is to match IP addresses to MAC addresses. The arp command allows us to view and modify the ARP table entries on our local computer. We can use this information to discover and correct any duplicate IP addresses or invalid entries that might cause connection issues.
One of the biggest reasons for connectivity issues is an improperly configured Domain Name System, or DNS. DNS is what's responsible for matching IP addresses to website or computer names. If we're running a web server, we absolutely need to make sure that our DNS entries are correct or no one will be able to connect. Even in an internal network, an improperly configured DNS can cause devices to be unable to connect to the server. We use the nslookup command in Windows or the dig command in Linux to view and modify our DNS settings.
Another big reason for connectivity issues is the IP configuration on the local computer. To view current information about our local system, we use the ipconfig command in Windows and the ifconfig command in Linux. These commands show us the current IP address, subnet mask, default gateway, and more. We can use this information to correct DHCP issues or static IP configurations.
Hping is a security tool that can be used to send ICMP packets as well as TCP, UDP, and RAW IP packets. Aside from simply checking connectivity, hping is used to analyze the target system to gather more information on it. The hping command is primarily designed for Linux but can be installed on Windows as well.
Netcat is an extremely powerful security tool that can basically do anything with TCP. Netcat is able to open up a TCP connection between two devices and send UDP packets between them, scan for open ports, and even listen in on connections between individual ports.
IP scanners are tools that allow us to scan our entire network in order to find all connected devices and their IP addresses. Depending on the program, these tools pull up basic-to-advanced information such as routes, hostnames, operating systems, and more.
Nmap is another powerful network monitoring tool. Nmap is a command line utility, but there is a GUI-based version called Zenmap. We can use nmap to scan an entire network or a specific IP address to discover all sorts of information like open ports, running services, operating system information, and beyond. One thing that makes nmap so powerful is that it can scan with different protocols on individual networks.
That'll wrap up this lesson. In this video, we covered a variety of tools and commands we can use to monitor and manage our network. It's important that we know each of these tools and how they're used so that we can pick the best tool for our situation.

6.2.2 Network Monitoring Facts

This lesson covers the following topics:

- Network monitoring overview
- Network monitoring tools

Network Monitoring Overview

The goal of monitoring is to keep track of conditions on the network, identify situations that might signal potential problems, pinpoint the source of problems, and locate areas of your network that might need to be upgraded or modified.

As you monitor your network, look for the top talkers and listeners.

- Top talkers are computers that send the most data, either from your network or into your network.
- Top listeners are hosts that are receiving most of the data by streaming or downloading large amounts of data from the internet.

It is important to know which computers are the big receivers and senders of information because it is a good way to tell if something is wrong on your network. An unauthorized system that is sending large amounts of data to locations outside of your network could be a sign of a data breach.

Monitoring also plays a vital role in endpoint hardening, helping to enforce and maintain the security measures put in place during the hardening process. Once devices are hardened, monitoring helps to ensure these conditions remain in place.

Security analysts can detect changes that weaken the hardened configuration through continuous monitoring. For instance, if a previously disabled port is detected as open or a service that was disabled is changed to enabled, monitoring tools can alert analysts of the change—which may indicate a breach.

Additionally, monitoring can provide valuable data for compliance and auditing purposes. Regular reports on the status of endpoint devices can verify that hardening baselines have been effectively deployed and maintained, supporting compliance with various regulations and industry standards.

Network Monitoring Tools

The below table lists some of the tools used to monitor the health of a network:

Tool	Description
ping	Ping is a command line tool that is used to perform a connection test between two network devices. Ping works by sending ICMP packets to a specified device on the network and waiting for a response. This shows if there is a connection issue or not. The syntax for the ping command is: **ping <target IP address or hostname>** The following switches are the more common switches that can be used to modify the ping command: • **-t** sends ICMP packets until manually stopped. • **-a** resolves addresses to hostnames. • **-n <*count*>** specifies the number of ICMP packets to send. Ping sends 4 packets by default • **-l <*size*>** specifies the packet size in bytes. ping sends 32-byte packets by default
tracert/traceroute	The tracert tool shows the path a packet takes to reach its destination. Every device the packet passes through is known as a hop. Use tracert to locate network devices that are down or causing latency issues. • tracert is the Windows version and sends ICMP packets. • traceroute is used in Linux and sends UDP packets.
pathping	The pathping Windows command line tool combines the tracert and ping tools. Use pathping to locate network devices that are down or causing latency issues.
netstat	Use the netstat command to display a variety of network statistics in both Windows and Linux, including: • Connections for different protocols • Open ports • Running programs Some of the common switches used to specify the information shown in Windows are:

	• **-a** displays all connections and listening ports. • **-b** displays the executable involved in creating each connection or listening port. • **-f** displays the FQDN for the foreign address if possible. • **-r** displays the routing table • **-p <*protocol*>** shows the connections for a specified protocol (TCP, UDP, TCPv6, UDPv6)
route	The route command is used in both Windows and Linux to show the routing table and to make manual changes to the table.
arp	The arp command is used in both Windows and Linux. ARP stands for Address Resolution Protocol and is used to match IP addresses to MAC addresses. The arp command displays, adds, and removes arp information from network devices. Some of the common switches used with the arp command are: • **-a** displays current ARP entries. • **inet_addr** specifies an internet address • **-d** deletes the host specified by inet_addr
nslookup/dig	The nslookup and dig commands are used to view and modify DNS settings. These tools can be used to look up DNS server information and also give IP addresses and domain names for a network server. • nslookup is used in Windows. • dig is used in Linux.
ipconfig/ifconfig	The ipconfig command (Windows) and the ifconfig command (Linux) are used to display the IP configuration on the local computer. Information such as the following can be shown using these commands: • Adapter name • Adapter MAC address • If DHCP is enabled or not • IPv6 address • IPv4 address • Subnet mask • IP lease information • Default gateway • DHCP server • DNS server
hping	Hping is a security tool that can check connectivity and also analyze the target to gather information. Hping can send ICMP, TCP. UDP, and RAW-IP packets. Hping is primarily designed for Linux but can be installed in Windows.

netcat	The netcat security tool can read and write data across both TCP and UDP network connections. It opens a TCP connection between two devices and can be used to send packets, scan for open ports, and listen in on connections to specific ports. You can download netcat from the internet.
IP scanners	IP scanners are special tools that allow a network administrator to scan the entire network to find all connected devices and their IP addresses. Advanced scans can also display information such as: • Routes • Hostnames • Operating systems
nmap	The nmap utility is a network security scanner. Use nmap to scan an entire network or specific IP addresses to discover all sorts of information such as: • Open ports • Running services • Operating system Nmap can use many different protocols and options depending on the network or device being scanned. Nmap is a command line tool, but a GUI version called Zenmap is available.

6.2.3 Performing Port and Ping Scans (Demo Video)

Transcript:

In this demo, we're going to look at Zenmap. Zenmap is the GUI equivalent of Nmap, which is a command line tool. We've downloaded and installed Zenmap for Windows. Even though Zenmap has a graphical user interface, we still use the familiar commands that Nmap uses. Other than that, the interface is quite simple.

First, you want to identify the target of your scan. In this case, it might be 192.168.2.1, a specific IP address, or maybe a range of IP addresses using CIDR notation. It's going to make sure that these first three octets stay the same, 192.168.2, then we'll try .1 and so on all the way up to .255 for the last octet.
To the right here, you can see what type of scan you want to perform. Right now, it's listed as an intense scan. To start out, I'm going do a lighter scan. You can see you have Intense Scan, plus UDP, plus all the TCP ports, without the ping, etc. You have the ping, the quick scan, the quick scan plus, and on and on. There's lots of different options here.
I'm going to start out with just a ping scan, which will simply see if the hosts in this IP address range respond to a ping. To do this, I'll simply select that option and then click on Scan. This will only take a few seconds to complete, as it's a very lightweight scan. You'll see here, as it goes on, that we'll start to populate the operating systems and the hosts associated with those IP addresses to the left.
In this case, we haven't done a deep scan. We don't know the operating systems on these hosts or other details yet. But we can see all the hosts that are in that range that responded to the ping scan. Maybe we come down and say, "Hey, this IP address is something of interest." Then we can go and do an intense scan with that specific host.
Before I do that, let me show you here that there's not really a lot we know about this host right now. We know its MAC address, its IP address, and whether it's currently up, but we don't know anything else.
You could do this intense scan with a full range of IP addresses. It'll just take a while. In the interest of time, I'll just do an intense scan on that single IP address.

This takes about four or five minutes to complete depending on the responsiveness of the server, so I'll stop the demo while it runs. But you can see that it's starting to discover open ports, different TCP ports, and it'll run through a variety of different scripts to try to discover information about this host.
When all is said and done, we should be able to see what type of operating system it is, even the version of the operating system, and other things that it uses as it discovers open ports. It can query those ports and ask certain questions in order to receive responses and determine the types of services that are running on this machine. We'll go ahead and pause now until that script is complete.
You can see that it's completed the scan. Let's look through some of these results first. We'll start at the top. It scanned for specific ports, and it started a service scan and gathered some information. It was looking for 20 specific services. It took about a minute to do that. Then it tried to detect the operating system and a variety of other services. It looks at the specific ports and tries to determine which programs are actually running. This server has DNS server. It's actually Microsoft's DNS server. Because of the types of queries that Nmap puts out, you can see the types of services that are actually running. We know that it's Microsoft DNS. We know Microsoft IIS, the web server, is running there, and we know the version of that web server and a variety of other services that are currently running.
Some remote procedure calls on certain ports that are open, and a variety of other information like the net BIOS names and the server versions. We can get quite a bit of detail using Nmap and the associated GUI. Let's click on 15. You can see that we have a summary of information for that. It's Windows, and it has the last boot time. It was back on May 9th. That's several months ago. Then we have the versions and other information. This is a great tool that you can use to quickly gather information about a network. It's easy to see why these tools make it so easy to attack and hack, or at least understand the composition of the machines within a network.
All right, that's it for this demo--just a quick demo of Nmap and the associated Zenmap GUI and how you can do ping scans and port scans using this tool.

6.2.4 Reconnaissance (Lesson Video)

Transcript:

Reconnaissance, which is also known as Footprinting, is the process of gathering information about a target before you begin an attack. There are two types of reconnaissance, passive and active. Passive reconnaissance is gathering information without directly interacting with users, and active reconnaissance is gathering information through direct interaction with users. In this lesson, we'll look at the tools associated with both types.

Before interacting with a target, a hacker should perform passive reconnaissance to gather as much information as possible without alerting the anyone. Gathering information through packet sniffing, eavesdropping, and Open-source intelligence, or OSINT, are the more common passive methods.
Packet sniffing is capturing data packets and analyzing them for useful information. If you're in a public place, you can gather these data packets and use a tool like Wireshark to see if you can find anything useful. Because you're not sending any data, you're essentially a bystander, so this is a passive reconnaissance method.
You need to scan for open wireless networks before you can start sniffing packets. Two common scanning methods are war driving and war flying.
War driving is driving or walking around with a wireless device to look for open vulnerable wireless networks. War flying is the doing the same thing using a drone or another unmanned aerial vehicle.
Eavesdropping is another way you can passively gather information. It can include listening to employee conversations, shoulder surfing, and dumpster diving. People often throw out documents with sensitive information and don't realize it. Holding your nose and sifting through dirty garbage doesn't require any special skills or tools. It's a really easy way for any hacker to get get really valuable information.
Open-source intelligence, or OSINT, is any data collected from publicly available sources. This includes doing web searches on search engines like Google, social media sites such as LinkedIn, and other sources to gather personal identifiable information, or PII.
After a hacker has gained as much information as possible using passive reconnaissance, they'll move on to active reconnaissance.
Active reconnaissance is the process of gathering information by interacting with the target. Because you are directly interacting with the target, you can be discovered.
During this phase, hackers gather additional information that they couldn't get with passive reconnaissance. This includes network information such as IP addresses, sub-domains, and DNS information. They'll also gather system

information such as operating systems, users and passwords, locations of servers, and organizational information that wasn't gathered during passive reconnaissance.
Doing all of this reconnaissance provides a wealth of knowledge on the target. The goal is to understand the target's security, narrow focus for potential attacks, identify vulnerabilities, and create a network map.
Sometimes the line between passive and active reconnaissance can get blurry. Just always remember that and any amount of direct interaction means you're using an active method.
There's a wealth of tools and resources hackers and penetration tester can use to gather all of this information.
One of the first tools a security professional should use in reconnaissance is the OSINT framework. This framework is a collection of resources and tools that are separated by common categories to easily gather all sorts of online information, making the initial reconnaissance process much easier.
theHarvester is a great passive reconnaissance tool that can gather lots of information from public sources, such as search engines, social media sites, and Shodan. There are some options in the program to brute force DNS information, which is active reconnaissance.
Shodan is a very popular search engine for internet-connected devices. Users can search for specific types of devices and locations. You can use this information to see if a target has any online devices without proper security and gather more information about them.
Dnsenum is a program that finds all DNS servers and entries for an organization. This helps you find information such as usernames, computer names, IP addresses, and more. Depending on the settings used, dnsenum can be either passive or active.
Curl and wget are two programs that let you download and upload files. You can use them in the active reconnaissance phase to download an entire website and view its code and resources offline. Because you're actively interacting with the target, these are active reconnaissance tools.
Scanless is a really cool tool for port scanning. It uses exploitation websites to perform port scans on your behalf. This means that you can scan your target anonymously. Even though you're not directly interacting with the target, this's still considered active reconnaissance.
Sn1per is a scanning tool that combines many tools into one program. You can use Sn1per to find all sorts of information, such as DNS, open ports, services, and vulnerabilities. This is definitely an active reconnaissance program.
Finally, we have Nessus. Nessus is a proprietary vulnerability scanner developed by Tenable. You can use Nessus to scan your target for known vulnerabilities that you can exploit to gain access to the target. Nessus has many different options available for scanning, which allows you to narrow your focus.
Using these tools allows you to gather all the information you could want about a person or organization. Network defenders can also use these same tools to discover what information is available to others and take necessary steps to remove or hide anything that shouldn't be accessed.
And that wraps up this lesson. In this video, we compared passive and active reconnaissance techniques. Passive means there is no direct interaction with users, and active means there is direct interaction. We also reviewed some the common tools used in the reconnaissance phase.

6.2.5 Perform Reconnaissance with Nmap (Demo Video)

Transcript:

One of the very first things you do as a penetration tester or ethical hacker is gather information. There are a lot of tools that can help you to do this. Most ethical hackers start with some very basic command line tools. In this demo, we're going to perform some reconnaissance using some of these basic tools. Our target is a domain set up by the Nmap Security Scanner Project and insecure.org. They've set up scanme.nmap.org to help people learn about Nmap and other tools. They authorize students to do scans as long as they're done within reason. Let's close the web browser.

To get started, we're going to go through some reconnaissance steps. First, I just want to make sure that the site is live, so let's ping it. I'll just type, 'ping scanme.nmap.org', press Enter, and I get a response. I'll hit Ctrl+C to stop the ping, and I can confirm that the site is live.
Now let's find the path to our target by performing a traceroute. On Windows, you do that by typing tracert. But we're on our Kali Linux system, so we need to type in 'traceroute scanme.nmap.org' and then Enter. I get over a dozen hops between me and the target. These, down here, aren't responding to the request, but my traceroute still continues. This is normal behavior. Let's clear our screen.

Now, I'm interested in finding the name server information for our target. I'm going to type in 'nslookup scanme.nmap.org' and press Enter. I get the IP addresses for both IP version 4 and version 6 along with the server names.
You can perform a whois search from the command line in Linux. To do that, I'll type in 'whois nmap.org' and press Enter. I'll scroll up, and here, I can see all the information about nmap that's available from whois. I'll clear the screen again.
Next, let's see what happens if we try to make a connection with the target. I know that port 80 is open because we were just at the website. Let's use Netcat to see if we can make a connection to the site on port 80. For that, I'll type in 'nc -v scanme.nmap.org 80' and press Enter. I get a response back, and it looks like I'm connected. Typically, if you want to make sure you're connected, you can type 'help', and you'll get a response. I do get a response. Not only that, I get some information about what sort of server they're running and even the version.
Let's try another. I'll use my up arrow to get back to our last command. I'll change the 80 to a 22 and press Enter. I get a connection again and see that they're running SSH, and I get the version number along with some other information. All right, let's move on. I have to press Ctrl+C to get out of this, and I'll also clear the screen.
You'll get a lot of use out of Nmap as a penetration tester and ethical hacker. We were just using Netcat to make connections to ports. But with 65,535 possible ports, that would take all day. We can use Nmap to do a scan to see what ports are open. To do a port scan, I'll type, 'nmap -sS scanme.nmap.org' and press Enter. The -sS parameter scans the 1,000 most common ports. I get some results back. I know what ports are open. I know the state of the port and what service is running. Now, if I want to continue working with Netcat, I have a list of ports I know are open.
Earlier, we did a ping to see if the target was alive. Let's repeat that now to get the IP address. I'll type 'ping scanme.nmap.org' and press Enter. I'll press Ctrl+C to stop the ping. Now, I want to do a ping sweep to see what other IPs might be associated with scanme.nmap.org. To do a ping sweep, I'll type in 'nmap -sn 45.33.32.1-255' and press Enter. Now I can go through this list and see which other IPs might be associated with scanme.nmap.org.
That's it for this demo. In this demo, we did some basic reconnaissance. First, we used some command line tools to find information about our target. Then we used Netcat to connect to open ports. Finally, we used nmap to perform a basic port scan and a ping sweep.

6.2.6 Reconnaissance Facts

Passive and active reconnaissance provides crucial information that helps penetration testers understand target systems and identify potential vulnerabilities to plan an attack effectively. A combination of active and passive reconnaissance techniques yields the most comprehensive information regarding the target environment during a penetration testing engagement.

This lesson covers the following topics:

- Passive reconnaissance
- Active reconnaissance
- Reconnaissance tools

Passive Reconnaissance

Passive reconnaissance involves gathering information on the target with no direct interaction with that target. Passive reconnaissance aims to gather intelligence on the target environment and identify potential vulnerabilities while generating minimal evidence of their actions. Valuable information can be gathered using passive reconnaissance. The following table shows some of the common passive reconnaissance methods:

Passive Reconnaissance Method	Description
Packet sniffing	Packet sniffing is the process of capturing data packets that are flowing across the network and analyzing them for important information. Modern networks should have

	good protection against network sniffing attacks, but there are occasional circumstances that allow an attacker to gather sensitive information from the data packets. Packet sniffing is most easily performed on open wireless networks. Because the attacker is not sending data or actively interacting with the target, this is considered passive reconnaissance. Scanning for open wireless networks needs to be done before packets can be sniffed. Two common methods are war driving and war flying. • War driving is driving around with a wireless device, looking for open, vulnerable wireless networks. • War flying uses drones or unmanned aerial vehicles to find open wireless networks.
Eavesdropping	Eavesdropping is the act of covertly listening in on a communication between other people. This can include: • Listening to employees' conversations without them knowing. • Shoulder surfing is an eavesdropping technique where the listener obtains passwords or other confidential information by looking over the shoulder of the target as the target logs on or types information. • Dumpster diving, which is also considered eavesdropping. When dumpster diving, the attacker goes through the trash to find important information that may have accidentally been thrown away.
Open-source intelligence (OSINT)	Open-source intelligence is any data that is collected from publicly available sources. The goal is to gather as much personal identifiable information (PII) as possible. This includes information found from resources such as: • Search engines (Google, Bing) • Social media (Facebook, LinkedIn) • Company websites (About sections of websites, company directories) • Media sources (news sites, interviews, articles) • Public government sources (property appraisal sites, public records)
Network traffic analysis	Network traffic analysis involves monitoring network traffic to identify patterns, devices, IP addresses, and potential vulnerabilities without actively generating traffic.

Active Reconnaissance

After an attacker has gained as much information as possible through passive reconnaissance, the next step is the active reconnaissance phase.

Active reconnaissance involves actively probing and interacting with target systems and networks to gather information. Active reconnaissance includes activities that generate network traffic by directly requesting information from target systems. Active reconnaissance aims to discover and obtain information about the target infrastructure, services, and potential vulnerabilities. Common techniques used in active reconnaissance include the following:

- Port scanning — scanning a target network to identify open ports and the services running on them.
- Service enumeration — interacting with identified services to gather information about their versions, configurations, and potential vulnerabilities.
- OS fingerprinting — attempting to identify the operating system running on target machines by analyzing network responses and behavior.
- DNS enumeration — gathering information about the target's DNS infrastructure, such as domain names, subdomains, and IP addresses.
- Web application crawling — exploring web applications to identify pages, directories, and potential vulnerabilities.

Performing reconnaissance provides the attacker with the information needed to perform a successful attack on the target. The goal is to know and understand the following information about the target:

- Security posture (this includes both network and physical security)
- How to narrow the focus for attack
- Potential vulnerabilities
- How to best create a network map

Reconnaissance Tools

There are many tools and resources available to assist in the reconnaissance phase. The table below covers some of the popular tools used:

Reconnaissance Tool	Description
OSINT framework	The OSINT framework is a collection of resources and tools that are separated by common categories. The OSINT Framework makes it easy to gather all sorts of information, making the initial reconnaissance process much more efficient. Documentation can be found at https://osintframework.com/
theHarvester	theHarvester is a passive reconnaissance tool that is used to gather information from a variety of public sources. The tool gathers emails, names, subdomains, IPs, and URLs using multiple public data sources. These sources include search engines, social media sites, and Shodan. theHarvester does have some options, such as brute-forcing DNS and taking screenshots, that would fall under active reconnaissance.
Shodan	Shodan is a popular search engine for internet-connected devices. Users are able to search for specific types of devices and locations. This information can be used to see if a target has any online devices without proper security.
Dnsenum	Dnsenum is a program that performs DNS enumeration and can find the DNS servers and entries for an organization. This information can help find other information such as usernames, computer names, IP addresses, and more.

Curl and wget	Curl and wget are two common command line programs that can be used to download or upload files. An example of using these tools is to download an entire website for offline analysis. Because these tools actively engage with the target, they are considered active reconnaissance tools.
Scanless	Scanless is used for port scanning. Instead of scanning ports from the hacker machine, scanless uses exploitation websites to perform port scans. This means the attacker is able to maintain anonymity while scanning the target.
Sn1per	Sn1per is an automated scanner that can be used to enumerate and scan for vulnerabilities. Sn1per combines the functions of many tools and can be used to find information such as DNS information, open ports, running services, and more.
Nessus	Nessus is a proprietary vulnerability scanner that was developed by Tenable. Nessus can be used to scan the target for any known vulnerabilities, which can be exploited to gain access to the target.

Selecting the right tool allows the attacker to gain the necessary information on the target. Network defenders can also use these tools to discover what information is out there and take the necessary steps to remove or hide anything that should not be available.

6.2.7 Practice Questions (Section Quiz)

q_netmon_dig_secp8

Which of the following tools can be used to view and modify DNS server information in Linux?

Answers:

- tracert
- ***dig**
- route
- netstat

Explanation:

The **dig** command is used to view and modify DNS settings. These tools can be used to look up DNS server information and give IP addresses and domain names for a network server.

The **tracert** command shows the path a packet takes to reach its destination. This is not the best tool for checking connectivity between two network devices.

The **route** command is used in both Windows and Linux to show the routing table and to make manual changes to it.

The **netstat** command is used to display a variety of network statistics in both Windows and Linux. This command is not used to look up DNS server information.

q_netmon_endpoint_hardening_secp8

In the context of network monitoring, which of the following is crucial for enforcing and maintaining the security measures put in place during the hardening process?

Answers:

- ***Endpoint hardening**
- Top talkers
- Packet routing
- IP scanning

Explanation:

Endpoint hardening is the correct answer. It refers to the process of securing individual devices on a network (endpoints) against threats. Network monitoring plays a vital role in endpoint hardening as it helps to enforce and maintain the security measures put in place during the hardening process. It can detect changes that weaken the hardened configuration and alert analysts of the change, which may indicate a breach.

Top talkers refers to the computers that send the most data, either from your network or into your network. While monitoring these can help identify potential network issues, it is not directly related to the enforcement and maintenance of security measures.

Packet routing refers to the path a packet takes to reach its destination. While this is an important aspect of network monitoring, it does not directly relate to the enforcement and maintenance of security measures.

IP scanning is a process that allows a network administrator to scan the entire network to find all connected devices and their IP addresses. While this is a useful tool in network monitoring, it does not directly enforce or maintain security measures.

q_netmon_hping_secp8

A tech startup, TechPioneers Inc., is looking to improve their network monitoring capabilities.

The company operates primarily on Linux-based systems and requires a tool that can not only check connectivity but also analyze targets to gather information. The tool should be capable of sending ICMP, TCP, UDP, and RAW-IP packets.

As a network administrator, which of the following tools would you recommend as the BEST solution for TechPioneers Inc.?

Answers:

- Ping
- Tracert/Traceroute
- ***Hping**
- Netstat

Explanation:

Hping is the best solution for TechPioneers Inc. Hping can check connectivity and also analyze the target to gather information. It can send ICMP, TCP, UDP, and RAW-IP packets. Although it is primarily designed for Linux, it can also be installed in Windows, making it a versatile choice for diverse network environments.

While ping is a useful tool for checking the connectivity between two network devices by sending ICMP packets and waiting for a response, it does not have the capability to analyze targets for gathering information or send TCP, UDP, and RAW-IP packets. Therefore, it is not the best solution for TechPioneers Inc.

Tracert/Traceroute is a tool that shows the path a packet takes to reach its destination, identifying each device the packet passes through. However, it does not have the capability to analyze targets for gathering information or send TCP, UDP, and RAW-IP packets. Therefore, it is not the best solution for TechPioneers Inc.

Netstat is a command-line tool that displays network statistics, including connections for different protocols, open ports, and running programs. However, it does not have the capability to analyze targets for gathering information or send TCP, UDP, and RAW-IP packets. Therefore, it is not the best solution for TechPioneers Inc.

q_netmon_nmap_secp8

You want to identify all devices on a network along with a list of open ports on those devices. You want the results displayed in a graphical diagram.

Which tool should you use?

Answers:

- Ping scanner
- OVAL
- Port scanner
- ***Network mapper**

Explanation:

A network mapper is a tool that can discover devices on a network and show those devices in a graphical representation. Network mappers typically use a ping scan to discover devices and a port scanner to identify open ports on those devices.

A ping scanner only identifies devices on a network, but does not probe for open ports.

A port scanner finds open ports, but it might not display devices in a graphical representation.

Open Vulnerability and Assessment Language (OVAL) is an international standard for testing, analyzing, and reporting the security vulnerabilities of a system.

q_netmon_ping_secp8

You need to check network connectivity from your computer to a remote computer.

Which of the following tools would be the BEST option to use?

Answers:

- tracert
- nmap
- ***ping**
- route

Explanation:

The **ping** command is used to perform a connection test between two network devices. It works by sending ICMP packets to a specified device on a network and waiting for a response. This shows if there is a connection issue or not.

The **tracert** command shows the path a packet takes to reach its destination. This is not the best tool to check for connectivity between two network devices.

The **nmap** utility is a network security scanner. Use **nmap** to scan an entire network or specific IP addresses to discover all sorts of information. This is not the best tool to check for connectivity between two network devices.

The **route** command is used in both Windows and Linux to show the routing table and to make manual changes to it.

q_recon_active_01_secp8

A company hires a security analyst to perform a penetration test on its network. During the process, the analyst plans to use various reconnaissance techniques to collect information about the target system.

In which of these reconnaissance methods does the security analyst directly interact with the target system?

Answers:

- ***Active**
- Passive
- Open-source intelligence (OSINT)
- Social engineering

Explanation:

Active reconnaissance involves probing and interacting with targeted systems to gather information about them.

Passive reconnaissance focuses on collecting publicly available data and passively observing network traffic without directly interacting with the target systems.

Open-source intelligence gathering is a passive reconnaissance technique that involves collecting publicly available information from various sources like search engines, social media, public databases, and websites.

Social engineering involves gathering information by manipulating people into revealing confidential information. It is part of passive reconnaissance since it does not involve direct interaction with the target systems.

q_recon_active_02_secp8

As a cybersecurity analyst, you are tasked with performing active reconnaissance on a potential client's network to identify vulnerabilities. You have already completed the passive reconnaissance phase.

Which of the following steps would you take next, and why?

Answers:

- Start by launching a denial-of-service (DoS) attack to test the network's resilience.
- ***Begin with port scanning to identify open ports and the services running on them.**
- Immediately report to the client that their network is secure based on the passive reconnaissance results.
- Use social engineering techniques to trick employees into revealing sensitive information.

Explanation:

Beginning with port scanning is the correct answer. Port scanning is a common technique used in active reconnaissance. It involves scanning a target network to identify open ports and the services running on them. This information can help identify potential vulnerabilities and areas for further investigation.

Launching a DoS attack during the reconnaissance phase is not ethical or legal. It could also potentially harm the client's network and disrupt their operations. Active reconnaissance should be non-disruptive and focused on gathering information, not causing harm.

While passive reconnaissance can provide valuable information, it does not provide a complete picture of the network's security. Active reconnaissance is necessary to further probe the network and identify potential vulnerabilities that were not apparent during the passive reconnaissance phase.

Social engineering techniques are not part of active reconnaissance. They involve manipulating individuals into revealing sensitive information, which is not ethical and could potentially be illegal. Active reconnaissance should focus on technical methods of gathering information about the network.

q_recon_nessus_secp8

You want to use a tool to scan a system for vulnerabilities, including open ports, running services, and missing patches.

Which tool should you use?

Answers:

- ***Nessus**
- LC4
- Wireshark
- OVAL

Explanation:

A vulnerability scanner is a software program that searches an application, computer, or network for weaknesses. These weaknesses could be things such as open ports, running applications or services, missing critical patches, default user accounts that have not been disabled, and default or blank passwords. Vulnerability scanning tools include Nessus, Retina Vulnerability Assessment Scanner, and Microsoft Baseline Security Analyzer (MBSA).

Wireshark is a protocol analyzer.

LC4 is a password-cracking tool that you can use to identify weak passwords.

Open Vulnerability and Assessment Language (OVAL) is an international standard for testing, analyzing, and reporting the security vulnerabilities of a system.

q_recon_osint_01_secp8

Gathering as much personally identifiable information (PII) on a target as possible is a goal of which reconnaissance method?

Answers:

- Active

- ***OSINT**
- Packet sniffing
- Passive

Explanation:

Open-source intelligence is any data that is collected from publicly available sources. The goal is to gather as much personally identifiable information (PII) as possible on the target.

Dumpster diving is when an attacker goes through the trash to find important information that may have accidentally been thrown away.

Active reconnaissance is the process of gathering information by interacting with the target in some manner.

Packet sniffing is the process of capturing data packets that are flowing across a network and analyzing them for important information.

q_recon_osint_02_secp8

An IT administrator has reviewed security policies and best practices on well-known IT bulletin boards. During this reading, the admin became aware of several new vulnerability exploits.

What type of information is this?

Answers:

- ***Open-source intelligence (OSINT)**
- Information-sharing organizations
- Cyber threat intelligence
- Threat feeds

Explanation:

Open-source intelligence (OSINT) describes collecting and analyzing publicly available information and using it to support decision-making. OSINT identifies vulnerabilities and threat information by gathering data from many sources such as blogs, forums, social media platforms, and even the dark web.

Information-sharing organizations are groups composed of businesses, government entities, and academic institutions. Each member of an organization benefits from sharing information with the group members.

Cyber threat intelligence is investigating, collecting, analyzing, and disseminating information about emerging threats and threat sources.

Threat feeds are signatures and pattern-matching rules supplied to analysis platforms as an automated feed. Companies respond swifter to emergent threats, as threat feeds provide real-time information.

q_recon_osint_03_secp8

A threat actor will use basic internet research to gather enough information about a local college to determine attack vectors.

Which of the following BEST describes passive reconnaissance that involves little physical work to accomplish the task?

Answers:

- ***OSINT**
- Bug bounty
- Service enumeration
- OS Footprinting

Explanation:

Open-source intelligence (OSINT) is any intelligence derived from publicly available information. The public information can include domains, IP address ranges, employees, and other data that will identify attack vectors.

Software vendors operate bug bounty programs where they give rewards for reporting vulnerabilities on their applications or systems.

Service enumeration is an active reconnaissance activity where people interact with identified services to gather information about their versions, configurations, and potential vulnerabilities.

Footprinting means scanning for hosts, IP ranges, and routes between networks to map out the structure of the target network.

q_recon_passive_01_secp8

Which type of reconnaissance is associated with dumpster diving?

Answers:

- Active
- OSINT
- Packet sniffing
- ***Passive**

Explanation:

Dumpster diving is when an attacker goes through the trash to find important information that may have accidentally been thrown away. Because there is no direct interaction with the target, dumpster diving is a form of passive reconnaissance.

Active reconnaissance is the process of gathering information by interacting with the target in some manner. Dumpster diving does not fall under this category.

Open-source intelligence (OSINT) is any data that is collected from publicly available sources. Dumpster diving does not fall under this category.

Packet sniffing is the process of capturing data packets that are flowing across a network and analyzing them for important information. Dumpster diving does not fall under this category.

q_recon_passive_02_secp8

Which passive reconnaissance tool is used to gather information from a variety of public sources?

Answers:

- ***theHarvester**
- Packet sniffing
- Shodan
- scanless

Explanation:

theHarvester is a passive reconnaissance tool that is used to gather information from a variety of public sources. This tool gathers emails, names, subdomains, IPs, and URLs using multiple public data sources. These include search engines, social media sites, and Shodan.

Packet sniffing is the process of capturing data packets that are flowing across a network and analyzing them for important information.

Shodan is a popular search engine for internet-connected devices. Users can search for specific types of devices and locations.

Use scanless for port scanning. Instead of an attacker scanning ports from their own machine, scanless uses exploitation websites to perform port scans on their behalf.

q_recon_recon_secp8

Which of the following tools can be used to see if a target has any online IoT devices without proper security?

Answers:

- theHarvester
- Packet sniffing
- ***Shodan**
- scanless

Explanation:

Shodan is a popular search engine for internet-connected devices. Users can search for specific types of devices and locations. This information can be used to see if a target has any online devices without proper security.

theHarvester is a passive reconnaissance tool that is used to gather information from a variety of public sources.

Packet sniffing is the process of capturing data packets that are flowing across the network and analyzing them for important information.

Use scanless for port scanning. Instead of the attacker scanning ports from their own machine, scanless uses exploitation websites to perform port scans on their behalf.

q_recon_theharvester_secp8

You are a cybersecurity analyst tasked with performing passive reconnaissance on a potential client's network.

You need to gather information from a variety of public sources including emails, names, subdomains, IPs, and URLs.

Which of the following tools would be most appropriate for this task?

Answers:

- Shodan
- ***theHarvester**
- Dnsenum
- OSINT framework

Explanation:

theHarvester is the correct answer. This tool is specifically designed for passive reconnaissance and can gather a wide range of information from various public sources including emails, names, subdomains, IPs, and URLs. It uses multiple public data sources including search engines, social media sites, and Shodan, making it the most appropriate tool for this task.

While Shodan is a powerful tool for finding internet-connected devices, it is not specifically designed for gathering a wide range of information from various public sources such as emails, names, subdomains, IPs, and URLs.

While Dnsenum is a useful tool for DNS enumeration, it is not as comprehensive as theHarvester for gathering a wide range of information from various public sources.

While the OSINT framework is a collection of resources and tools for open-source intelligence, it does not have the specific functionality of gathering emails, names, subdomains, IPs, and URLs from various public sources like theHarvester does.

q_recon_war_driving_secp8

Which of the following is known as the process of walking around an office building with an 802.11 signal detector.

Answers:

- War dialing
- Daemon dialing
- Driver signing
- ***War driving**

Explanation:

War driving is the act of searching for wireless networks (802.11) using a signal detector or a network client (such as a PDA or notebook). While the phrase war driving originated from the action of driving around a city searching for wireless networks, the name currently applies to any method of searching for wireless networks, including walking around.

War dialing and daemon dialing are both the act of dialing phone numbers in search of an answering modem. Often, war/daemon dialing calls all of the phone numbers in an area code or a prefix range in search of active modems.

Driver signing is a method of signing device drivers in an attempt to verify the source and quality of installed drivers. However, signing a device driver only indicates its source. Signing does not guarantee the reliability, stability, quality, or compatibility of a device driver.

6.3 Intrusion Detection

As you study this section, answer the following questions:

- What is the difference between an IDS and IPS?
- Which component gathers data from source devices?

- Why is a false negative the worst possible action by an IDS?
- Which detection method causes more false negatives?
- What is the difference between a host-based and network-based IDS/IPS implementation method?

In this section, you will learn to:

- Implement intrusion detection and prevention.

The key terms for this section include:

Term	Definition
Intrusion detection system (IDS)	Device or software that monitors, logs, and detects security breaches, but takes no action to stop or prevent the attack.
Intrusion prevention system (IPS)	Device that monitors, logs, detects, and can also react to stop or prevent security breaches.
Sensor	IDS component that passes data from the source to the analyzer.
Engine	IDS component that analyzes sensor data and events, generates alerts, and logs all activity.
Signature-based detection	Also referred to as pattern matching, dictionary recognition, or misuse-detection (MD-IDS). This detection method looks for patterns in network traffic and compares them to known attack patterns called signatures.
Heuristic-based detection	Also referred to as behavior, anomaly, or statistical-based detection. This detection method first defines a baseline of normal network traffic and then monitors traffic looking for anything that falls outside that baseline.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Hardening techniques
	  ○ Host-based intrusion prevention system (HIPS)
	
	3.2 Given a scenario, apply security principles to secure enterprise infrastructure.
	• Infrastructure considerations
	  ○ Device attribute
	  ▪ Inline vs. tap/monitor

	○ Network appliances ▪ Intrusion prevention system (IPS)/intrusion detection system (IDS) ▪ Sensors 4.3 Explain various activities associated with vulnerability management. • Analysis ○ False positive ○ False negative 4.5 Given a scenario, modify enterprise capabilities to enhance security. • IDS/IPS ○ Trends ○ Signatures • User behavior analytics 5.6 Given a scenario, implement security awareness practices. • Anomalous behavior recognition ○ Unexpected
TestOut Security Pro	5.2 Assessment techniques 5.2.1 Implement intrusion detection

6.3.1 Intrusion Detection (Lesson Video)

Transcript:

The first step in defending our network against unauthorized access is knowing if someone is attempting to gain access. To assist network defenders in this task, we can implement network intrusion detection or network intrusion prevention systems. In this lesson, we'll look at the differences, detection methods, and implementation of these systems.

To better understand the differences between an intrusion detection and intrusion prevention system, lets take a look at what each of them do.
An intrusion detection system, or IDS, is used to monitor and detect unauthorized access. It doesn't try to stop an attack or prevent traffic; it just monitors data packets and creates logs. An IDS can be either a software program or a hardware device.
Similar to an IDS is an intrusion prevention system, or IPS. Not only does the IPS monitor network traffic, it also works to block unauthorized and suspicious packets from entering the network.
No solution is perfect, and that's why we should always implement security-in-depth and use both detection and prevention systems. Even with an IPS, malicious packets may still find their way into the network. The IDS can serve as a backup and alert us that something suspicious may be going on and record it for us. This information will be extremely valuable after the attack so that we can figure out exactly what happened and what the damage is.
Now that we understand what an IDS and IPS are, lets take a look at how they work.
Both IDS and IPS systems monitor data packets for suspicious traffic. They do vary slightly in how they accomplish this, but both systems use generally the same methods. For simplicity's sake we'll talk about the IDS systems, but understand that all of this can apply to both unless otherwise specified.

This first and most common method is signature-based detection. This type of IDS is sometimes called a misuse-detection IDS. It attempts to match the known traffic patterns of specific network attacks. Similar to how viruses have a unique fingerprint that antivirus programs use to detect them, malicious packets also have a unique fingerprint that the IDS can use to identify them with.
The other detection method is heuristic, or behavior-based detection, sometimes referred to as statistical anomaly-based detection. With this type of detection, you set up a baseline of regular network activity and the IDS determines if activity is outside the norm. The biggest issue with this method is that a larger number of false positives are likely to be generated.
No matter the detection method being used, an IDS can either be a host-based or network-based IDS.
A host-based IDS is software installed on the host system. Host-based IDSs are better at detecting attacks unique to the services and applications on that system than network-based IDSs. A host-based IDS also monitors local files for suspicious activity, which is something a network-based IDS can't do. The problem with host-based IDSs is that software must be installed on each system you want to protect. This can lead to excessive administrative effort. Also, if a system is compromised, the log reports on that system become unreliable because the attacker may have modified the log files.
The other option is to install the IDS on the network itself. This is a physical device that monitors network traffic in real time. We have two options on where to install a network IDS: inline or out of band.
Because an IDS is a passive device that simply monitors packets, they're generally installed as out-of-band devices. All this means is that the device is installed outside the flow of traffic. It's usually connected with a network tap, like a switch, which allows it to monitor traffic without being in the way.
An IPS on the other hand is an active device that's usually installed inline, or in the flow of traffic. All data goes through the IPS, where it's checked out and either allowed to pass through or stopped.
Regardless of whether the device is inline or out of band, it can be installed in front of or behind the firewall. Generally, an IDS or IPS should be installed behind the firewall as this uses less resources.
That'll wrap things up for this lesson. In this lesson, we looked at intrusion detection and intrusion prevention systems. The biggest difference between the two is that an IDS is a passive monitoring device and an IPS is an active prevention device. We also looked at the common detection methods that both use, which are signature-based detection and heuristic, or behavior-based, detection. Finally, we looked at how to install these devices on a network. IDSs should be installed out of band and IPSs should be inline.

6.3.2 IDS Facts

The first step in defending a network against unauthorized access is knowing that someone is gaining access. An intrusion detection system (IDS) is a special network device that can detect attacks and suspicious activity. An active IDS is known as an intrusion prevention system (IPS).

This lesson covers the following topics:

- Differences between an IDS and IPS
- Detection methods
- Device implementation
- Trend Analysis

Differences between an IDS and IPS

The table below shows the differences between an IDS and an IPS:

IDS	IPS
A passive IDS monitors, logs, and detects security breaches but takes no action to stop or prevent the attack. A passive IDS:	An active IDS, also called an IPS, performs the functions of an IDS but can

• Can send an alert, but this requires the security administrator to interpret the degree of the threat and respond accordingly. • Cannot be detected on the network because it takes no detectable actions.	also react when security breaches occur. An IPS: • Can automate responses to malicious or suspicious traffic. • Can terminate sessions (using the **TCP-RST** command) or restart other processes on the system. • Performs behaviors that can be seen by anyone watching the network. Usually, these actions are necessary to block malicious activities or discover the identity of an intruder. Updating filters and performing reverse lookups are common behaviors of an

	active IDS.

Using both of these devices in a network provides the best network detection and protection. If a malicious packet makes it past the IPS, the IDS serves as a backup and alerts the security operations team.

The IDS also records and logs everything (this can be viewed in the follow-up).

The steps an IDS/IPS takes when monitoring traffic are:

- A sensor passes data from the source to the analyzer.
- The engine, or analyzer, analyzes the sensor data and events, generates alerts, and logs all activity. An *alert* is a message indicating an event of interest (such as a possible attack)
- The IDS/IPS labels traffic based on its interpretation of whether or not the traffic poses a threat, as described in the following table.

State	Description
Positive	A positive traffic assessment means that the system detected an attack, the appropriate alarms and notifications were generated, or the correct actions were performed to prevent or stop the attack.
False positive	A false positive traffic assessment means that the system identified harmless traffic as offensive and generated an alarm or stopped the traffic.
Negative	A negative traffic assessment means that the system deemed the traffic harmless and let it pass.
False negative	A false negative traffic assessment means that harmful traffic was allowed to pass without any alerts being generated or any actions being taken to prevent or stop it. This is the worst possible scenario.

Detection Methods

Both systems monitor data packets for malicious or unauthorized traffic. The table below shows the different methods they can use to distinguish attacks and threats from normal traffic:

Detection Method	Description
Signature-based	Signature-based detection also referred to as pattern matching, dictionary recognition, or misuse-detection (MD-IDS), looks for patterns in network traffic and compares them to known attack patterns called signatures. Similar to how viruses have a unique fingerprint that antivirus programs use to detect their presence, malicious packets have a unique fingerprint that the IDS can use to do the same. These fingerprints are referred to as signatures. • Signatures are written and updated by the IDS vendors.

	• Signature-based detection cannot detect unknown attacks; it can only detect attacks identified by published signature files. For this reason, it is important to update signature files on a regular basis. Commercial software requires a paid-for subscription to obtain the updates. It is important to configure software to update only from valid repositories, ideally using a secure connection method such as HTTPS. • Signature-based detection usually causes more false negatives than heuristic-based detection.
Heuristic-based	Heuristic-based detection, also referred to as behavior, anomaly, or statistical-based detection, first defines a baseline of normal network traffic and then monitors it. It looks for anything that falls outside that baseline. • Clipping levels, or thresholds, are defined and used to identify deviations from the baseline. • When the threshold is reached, an alert is generated, or action is taken. • Heuristic-based systems can recognize and respond to some unknown attacks (attacks that do not have a corresponding signature file). • This detection method usually causes more false positives than signature-based detection.

Device Implementation

An IDS/IPS can be implemented as a host-based or network-based device. The table below describes each implementation:

Implementation Method	Description
Host-based	A *host-based* IDS (HIDS) is a program installed on the host system itself that monitors all traffic coming into the host. A host-based IDS: • Is used to detect attacks that are unique to the services and applications on that system. It can monitor application activity and modifications as well as local system files, logon audit files, and kernel audit files. • Is typically unaware of other devices on the network but can be detected and could be the target of an attack itself. • May rely on auditing and logging capabilities of the operating system. • Can analyze encrypted traffic (because services running on the host decrypt the traffic). Antivirus software is the most common form of host-based IDS. One issue with host-based IDSs is that the software must be installed and configured on each system being protected. This can lead to excessive administrative effort. Also, if the host system is compromised, the log reports on that system become unreliable because the attacker may have modified the log files.

Network-based	A *network-based IDS* (NIDS) is a dedicated device installed on the network. It analyzes all traffic on the network in real-time. There are two options when installing an NIDS:
	• The first option is to install the NIDS out of the band. This means it is installed outside the flow of traffic.
	  ○ The IDS is usually connected with a network tap, such as a switch. This allows it to monitor network traffic without being in the way.
	• The other option is to install the NIDS as an *inline* device. This means it is installed in the flow of traffic, and all traffic goes through the NIDS. It is then analyzed and either allowed to continue or is stopped.
	
	Some other things to be aware of when implementing an NIDS are:
	• An NIDS is typically unaware of individual hosts on the network. It cannot be detected by attacking systems.
	• An NIDS is particularly well suited for detecting and blocking port scanning and DoS attacks.
	• An NIDS is unable to analyze encrypted traffic.
	• An NIDS should be placed at all critical junctions within a network, including backbones and critical choke points, such as:
	  ○ Inside the DMZ.
	  ○ Between the firewall and the internal LAN.
	  ○ Near any critical information assets.
	  ○ If using a switch on the network, the NIDS must be placed on a special port called a *spanning* or *diagnostic* port that directly connects to the backbone of the switch. This way, the NIDS can see all traffic on that segment.
	• A control center should be set up to receive all IDS data. This is where all decision-making should take place regarding NIDS communications.
	• An application-aware NIDS can analyze network packets to detect malicious payloads targeted at Application layer services (such as a web server).

Trend Analysis

Trend analysis is a critical aspect of managing intrusion detection systems (IDS) and intrusion prevention systems (IPS) as it aids in understanding an environment over time, helping to identify patterns, anomalies, and potential threats. Security analysts can identify patterns and trends that indicate ongoing or growing threats by tracking events and alerts. For example, an increase in alerts related to a specific attack may suggest that a network is being targeted for attack or that a vulnerability is being actively exploited. Trending can also help in tuning IDS/IPS systems. Over time, security analysts can identify false positives or unnecessary alerts that appear frequently. These alerts can be tuned down so analysts can focus on more important alerts.

Trending data can contribute to operational security strategies by identifying common threats and frequently targeted systems. This approach highlights areas of weakness that need attention, either through changes in security policy or investment in additional security tools and training.

6.3.3 Implement Intrusion Detection and Prevention (Demo Video)

Transcript:

Intrusion detection and prevention is an important task that is required to protect todays networks. In this demonstration, we'll configured intrusion detection and prevention using Snort on a pfSense security appliance.

There are several products available that can do both intrusion detection and intrusion prevention. Some cost money and some are free.
One of the most popular products is Snort. Snort is a free, open source network intrusion detection system (IDS) and intrusion prevention system (IPS).
Although Snort is open source and free, it does have some paid plans that you can subscribe to and receive updates to rules faster. In this demo, we'll configure Snort on our pfSense security appliance.
Before we configure Snort, let's quickly look at the website. As we scroll down, you see that you can manually download Snort for various operating systems and distributions. Snort is available for Fedora, Centos, FreeBSD and Windows.
In the second step, you must get what is called an Oinkcode. We will do that in a minute but let's keep scrolling down this page. Step 3 is to get updates. We'll cover that later in the demo.
Now as a review what we just talked about, Snort is an open source intrusion prevention system capable of real-time traffic analysis and packet logging.
Snort is the most widely deployed IPS in the world. There have been over 5 million downloads and over 600 thousand registered users.
To use Snort, you must first get an Oinkcode. We'll go to our account. Here is a link that says Oinkcode. We already have a code generated, so we can proceed. We will use the Oinkcode later.
Now let's go to pfSense and install Snort.
We've already logged into pfSense and we're on the Dashboard. Installing Snort on pfSense is quite easy. It is done with the Package Manager that is located under the System tab. Once on the Package Manager page, you click Available Packages.
Now let's do a search for â€˜Snort'. We have one result, so we'll come down and click the Install link. Now you can make sure you have the right package. We do; so let's click Confirm to start the installation.
Now we wait for a few minutes while Snort is installed. When it does, the color changes from red to green and it says that the installation successfully completed. Down here on the bottom, a message says "Success".
Now that Snort is installed, we need to set it up. Let's do that by going to Services and then down to Snort. Be aware that this menu item for Snort was not there until we installed it. If you don't have this menu, Snort probably is not installed.
We will start by going to Global Settings. Under Snort Subscriber Rules, we check the Enable Snort VRT box. VRT is an acronym for Vulnerability Research Team.
Below that, next to Snort Oinkmaster Code, we paste in the code. Remember, we got the Oinkcode from the Snort website that we were at in the beginning of the demo. We copied the code to the clipboard and will use a keyboard shortcut to paste it in here.
Next, we Enable Snort GPLv2. The Community Snort Rules fall under the GNU General Public License Version 2, which encourages the development and distribution of open source software. This ruleset is 30 days behind the Snort Subscriber Rule Set. It does not contain zero-day threats under the limited provision of the Snort Subscriber Rule Set License.
Now let's check Enable ET Open (ET is the acronym for Emerging Threats). This downloads the Emerging Threats Open rules. The ET Open Ruleset is an anti-malware IDS/IPS ruleset that enables users with cost constraints to enhance their existing network-based malware detection.
We do not pay for the Emerging Threat Pro rules, so we won't check that box.
Let's skip down here under Sourcefire OpenAppID Detectors. Let's check the Enable OpenAppID box. Below that, check the box next to Enable RULES Open AppID. OpenAppID is an application-focused detection language and processing module for Snort. When you use OpenAppID with pfSense, the system can successfully detect (if configured to do so) and block over 2600 different services like Facebook, Netflix, Twitter, and Reddit.
For our Rules Update Settings, we set the Update Interval to 1 Day. For our Update time, we set it to 2:00 AM. Let's check the Hide Deprecated Rules Categories box. This removes old and outdated rules.
Under General Settings > Remove Blocked hosts Interval, we change that to 1 Hour. Now you might think that we should block hosts forever if they are malicious, but the problem is that often these are coming from spoofed IP addresses or from addresses that may be used by legitimate users very soon. So, we will block only for 1 hour.
We'll check the box for Startup/Shutdown Logging. We want to know who and when Snort is being started and stopped by.
That wraps up everything for this page. We'll click Save.
Now that we have the rules figured out, we need to assign these rules to the WAN interface. We'll go to the Snort Interfaces tab. Then come down here to the right and click the Add link.

Under General Settings, we want to make sure that Enable Interface is checked. We see that it is. Make sure you have WAN selected under Interface. This is the traffic we want to inspect. For Description, we enter WAN. We want to keep it obvious and simple here.
Under Alert Settings, we check the box to Send Alerts to System Log. This will send alerts to the firewall log.
We check the Block Offenders box. If an offender creates a Snort alert, they will be blocked. For the IP to Block, we block the Source IP address.
There is nothing more we want to do here. We'll come down and click Save.
Now we want to check the box to enable the WAN interface. Under Snort Status, click Runto start Snort.
That's it for this demo. In this demo we configured intrusion detection and prevention using Snort on our pfSense security appliance.

6.3.4 Implement Intrusion Prevention (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. In an effort to protect your network against security threats and hackers, you have added Snort to pfSense. With Snort already installed, you need to configure rules and settings and then assign Snort to the desired interface.

In this lab, your task is to use pfSense's Snort to complete the following:

Sign in to pfSense using the following:

- Username: **admin**
- Password: **P@ssw0rd** (zero)

- Enable the downloading of the following:
  - Snort free registered User rules
    - Oinkmaster Code: **359d00c0e75a37a4dbd70757745c5c5dg85aa**
  - Snort GPLv2 Community rules
  - Emerging Threats Open rules
  - Sourcefire OpenAppID detectors
  - APPID Open rules
- Configure rule updates to happen once a day at 1:00 a.m.
  - Hide any deprecated rules.
- Block offending hosts for 1 hour.
- Send all alerts to the system log when the Snort starts and stops.
- Assign Snort to the WAN interface using a description of **WANSnort** .
  - Include:
    - Sending alerts to the system log
    - Automatically blocking hosts that generate a Snort alert
- Start Snort on the WAN interface.

Explanation

Complete this lab as follows:

1. Sign in to the pfSense management console.
   a. In the Username field, enter **admin** .
   b. In the Password field, enter **P@ssw0rd** (zero).
   c. Select **SIGN IN** or press **Enter** .
2. Access the Snort Global Settings.
   a. From the pfSense menu bar, select **Services** > **Snort** .

b. Under the Services breadcrumb, select **Global Settings** .

3. Configure the required rules to be downloaded.
    a. Select **Enable Snort VRT** .
    b. In the *Sort Oinkmaster Code* field, enter **359d00c0e75a37a4dbd70757745c5c5dg85aa** . You can copy and paste this from the scenario.
    c. Select **Enable Snort GPLv2** .
    d. Select **Enable ET Open** .
4. Configure the Sourcefire OpenAppID Detectors to be downloaded.
    a. Under Sourcefire OpenAppID Detectors, select **Enable OpenAppID** .
    b. Select **Enable RULES OpenAppID** .
5. Configure when and how often the rules will be updated.
    a. Under Rules Update Settings, use the Update Interval drop-down menu to select **1 Day** .
    b. For Update Start Time, change to **01:00** .
    c. Select **Hide Deprecated Rules Categories** .
6. Configure Snort General Settings.
    a. Under General Settings, use the Remove Blocked Hosts Interval drop-down menu to select **1 HOUR** .
    b. Select **Startup/Shutdown Logging** .
    c. Select **Save** .
7. Configure the Snort Interface settings for the WAN interface.
    a. Under the Services breadcrumb, select **Snort Interfaces** and then select **Add** .
    b. Under General Settings, make sure **Enable interface** is selected.
    c. For Interface, use the drop-down menu to select **WAN (CorpNet_pfSense_L port 1)** .
    d. For Description, use **WANSnort** .
    e. Under Alert Settings, select **Send Alerts to System Log** .
    f. Select **Block Offenders** .
    g. Scroll to the bottom and select **Save** .
8. Start Snort on the WAN interface.
    a. Under the Snort Status column, select the **arrow** .
    b. Wait for a checkmark to appear, indicating that Snort was started successfully.

6.3.5 Practice Questions (Section Quiz)

q_ids_active_passive_secp8

An organization in the education sector plans to implement an Intrusion Detection System (IDS) with strategically placed sensors as part of its network infrastructure redesign to improve enterprise security.

The IT manager also considers implementing an Intrusion Prevention System (IPS) and deliberates over active versus passive modes.

Which statement BEST describes the implications of choosing an active IPS over a passive IDS?

Answers:

- ***An active IPS can block malicious traffic but may introduce latency, while a passive IDS avoids latency.**
- An active IPS can only detect and alert on malicious traffic, while a passive IDS can block such traffic.
- Both an active IPS and a passive IDS can block malicious traffic, but the IPS causes more latency.
- Both an active IPS and a passive IDS can only detect and alert on malicious traffic and cause latency.

Explanation:

An active IPS can block malicious traffic but may introduce some latency into the network traffic. A passive IDS cannot block traffic. It can only detect and alert malicious traffic and does not cause latency.

It is inaccurate to state that an active IPS can only detect malicious traffic, and a passive IDS can block such traffic. In reality, it's the other way around.

Active IDS can block malicious traffic, while passive IDS cannot block malicious traffic.

An active IPS can do more than detect and alert malicious traffic, it can also block such traffic.

q_ids_analysis_engine_secp8

A small company recently installed an intrusion detection system (IDS).

What is the purpose of the analysis engine in the IDS?

Answers:

- To capture network traffic
- ***To interpret and scan captured traffic for suspicious activity**
- To generate incident reports for security analysts
- To update the signatures and rules for attack patterns

Explanation:

An IDS's analysis engine scans and interprets the captured traffic to identify suspicious activity, classify events (e.g., ignore, log, alert, or block), and generate incident reports.

The sensor, not the analysis engine, is responsible for capturing network traffic in an IDS.

The analysis engine is not designed to generate incident reports.

The analysis engine is not designed to update the signatures and rules for attack patterns.

q_ids_anomaly_01_secp8

You are concerned about protecting your network from network-based attacks on the internet. Specifically, you are concerned about attacks that have not yet been identified or that do not have prescribed protections.

Which type of device should you use?

Answers:

- Signature-based IDS
- ***Anomaly-based IDS**
- Antivirus scanner
- Network-based firewall
- Host-based firewall

Explanation:

An anomaly-based intrusion detection system (IDS) can recognize and respond to some unknown attacks. Signature recognition, also referred to as pattern matching or dictionary recognition, looks for patterns in network traffic and compares them to known attack patterns called signatures. Signature-based recognition cannot detect unknown attacks. This system can only detect attacks identified by published signature files.

Antivirus software is a form of signature-based IDS. A network-based firewall filters packets for a network, while a host-based firewall filters packets for a host. Firewalls are typically configured using access control lists that identify specific traffic as allowed or denied.

q_ids_anomaly_02_secp8

Which intrusion detection method involves the analysis engine trained to recognize baseline "normal" traffic and generates an incident when it detects deviations from this baseline?

Answers:

- Signature-based detection
- ***Behavioral- and anomaly-based detection**
- Trend analysis
- Network traffic analysis (NTA)

Explanation:

Behavioral- and anomaly-based detection involves training the engine to recognize any activity that deviates from this baseline (outside a defined tolerance level) to generate an incident. It helps identify zero-day attacks, insider threats, and other malicious activities without a specific signature.

Signature-based detection relies on a database of attack patterns to identify malicious traffic. Effective against known threats, but signature-based detection may miss new or evolving attacks.

Trend analysis aids in understanding an environment over time, helping to identify patterns, anomalies, and potential threats. However, it does not involve the analysis engine.

Network traffic analysis (NTA) is a method of monitoring network availability and activity to identify anomalies, including security and operational issues. However, it does not involve the analysis engine.

q_ids_false_neg_secp8

Which of the following describes the worst possible action by an IDS?

Answers:

- The system detected a valid attack and the appropriate alarms and notifications were generated.
- The system correctly deemed harmless traffic as inoffensive and let it pass.
- The system identified harmless traffic as offensive and generated an alarm.
- ***The system identified harmful traffic as harmless and allowed it to pass without generating any alerts.**

Explanation:

The worst possible action an IDS can perform is identifying harmful traffic as harmless and allowing it to pass without generating any alerts. This condition is known as a false negative.

Positive traffic assessment means that the system detected a valid attack and the appropriate alarms and notifications were generated. Negative traffic assessment means that the system correctly deemed harmless traffic as inoffensive and let it pass. False positive traffic assessment means that the system identified harmless traffic as offensive and triggered an alarm.

q_ids_false_pos_01_secp8

You have configured an NIDS to monitor network traffic.

Which of the following describes harmless traffic that has been identified as a potential attack by the NIDS device?

Answers:

- ***False positive**
- False negative
- Negative
- Positive

Explanation:

A false positive traffic assessment means that the system identified harmless traffic as offensive and generated an alarm or stopped the traffic.

A negative traffic assessment means that the system deemed the traffic harmless and let it pass.

A false negative traffic assessment means that harmful traffic was allowed to pass without any alerts being generated or any actions being taken to prevent or stop it. This is the worst possible action by an IDS.

A positive traffic assessment means that the system detected an attack and the appropriate alarms and notifications were generated or the correct actions were performed to prevent or stop the attack.

q_ids_false_pos_02_secp8

Which of the following describes a false positive when using an IPS device?

Answers:

- ***Legitimate traffic being flagged as malicious.**
- Malicious traffic not being identified.
- Malicious traffic masquerading as legitimate traffic.
- The source address matching the destination address.
- The source address identifying a non-existent host.

Explanation:

On an intrusion prevention system (IPS), a positive match occurs when traffic matches the signature that identifies malicious traffic. A false positive occurs when legitimate traffic is identified as malicious traffic. This situation is undesirable, as it often results in legitimate traffic being rejected. Good IPS signature files result in low false positive rates.

A false negative occurs when malicious traffic is not identified and is, therefore, allowed.

Spoofing is the technique of falsifying the source address in a packet.

q_ids_firewall_rules_secp8

An organization's cybersecurity team has recently set up a new firewall and intrusion detection system (IDS) to strengthen the security of its enterprise infrastructure. The IDS, however, has sent a high number of false positive alerts, which hampers efficient threat monitoring.

The team believes adjusting the firewall rules will decrease these false positives without weakening the network's security.

What strategy should the cybersecurity team implement to fine-tune their firewall rules to reduce the IDS's false positives, ensuring a robust security infrastructure?

Answers:

- Disable firewall rules for less critical services.
- ***Establish firewall rules based on service-specific threat intelligence.**
- Increase the IDS's sensitivity level.
- Allow all traffic through the firewall.

Explanation:

By basing firewall rules on specific threat intelligence, the team can more accurately filter out potential threats while reducing false positives and maintaining security.

Disabling firewall rules for less critical services exposes the network to threats targeting those services, leading to possible security breaches.

Increasing the IDS's sensitivity level could increase false positives rather than decrease them, adding more noise to threat detection.

Allowing all traffic through the firewall would be highly risky as it would expose the network to potential threats, making it easier for malicious entities to gain access to the network.

q_ids_host_01_secp8

As a security precaution, you have implemented IPsec that is used between any two devices on your network. IPsec provides encryption for traffic between devices.

You would like to implement a solution that can scan the contents of the encrypted traffic to prevent any malicious attacks.

Which solution should you implement?

Answers:

- ***Host-based IDS**
- Network-based IDS
- VPN concentrator
- Port scanner
- Protocol analyzer

Explanation:

A host-based IDS is installed on a single host and monitors all traffic coming into the host. A host-based IDS can analyze encrypted traffic because the host operating system decrypts that traffic as it is received.

A network-based IDS is a dedicated device installed on the network. It analyzes all traffic on the network. It cannot analyze encrypted traffic because the packet contents are encrypted so that only the recipient can read the packet contents.

A protocol analyzer examines packets on the network, but it cannot look at the contents of encrypted packets.

A port scanner probes a device to identify open protocol ports.

A VPN concentrator is a device used to establish remote access VPN connections.

q_ids_host_02_secp8

What do host-based intrusion detection systems often rely upon to perform detection activities?

Answers:

- Network traffic
- ***Auditing capabilities**
- External sensors
- Remote monitoring tools

Explanation:

A host-based IDS often relies upon the host system's auditing capabilities to perform detection activities. The host-based IDS uses the logs of the local system to search for attack or intrusion activities.

Host-based IDSs do not analyze network traffic, use external sensors, or rely upon remote monitoring tools.

q_ids_host_03_secp8

What is the MOST common form of host-based IDS that employs signature or pattern-matching detection methods?

Answers:

- ***Antivirus software**
- Firewalls
- Honeypots
- Motion detectors

Explanation:

Antivirus software using signatures is the most commonly deployed form of a host-based IDS.

Firewalls are primarily designed to secure the network from cyberattacks. They check all incoming and outgoing traffic on a network according to a defined set of rules, not signature or pattern-matching detection methods.

While honeypots help to refine an organization's intrusion detection system (IDS) and threat response, they do so by mimicking a target for hackers, and uses their intrusion attempts to gain information about cybercriminals and the way they are operating or to distract them from other targets.

Motion detectors are hardware that sense physical movements to perform tasks such as helping to turn on outside lights. They do not employ signature or pattern-matching detection methods.

q_ids_ids_01_secp8

Which security mechanism can be used to detect attacks that originate on the internet or from within an internal trusted subnet?

Answers:

- Firewall
- ***IDS**
- Security alarm
- Biometric system

Explanation:

An IDS is a security mechanism that can be used to detect attacks that originate on the internet or from within an internal trusted subnet.

A firewall is only able to filter traffic crossing through its interfaces. So, firewalls are unlikely to detect attacks within trusted areas that are not filtered.

A security alarm is a form of physical intrusion detection. In other words, it is not suited for detecting electronic or online attacks.

Biometric systems are a form of an authentication mechanism, not a detection mechanism.

q_ids_ids_02_secp8

Which of the following is a security service that monitors network traffic in real time or reviews the audit logs on servers looking for security violations?

Answers:

- Firewall
- Switch
- Padded cell
- ***IDS**

Explanation:

An IDS (intrusion detection system) is a security service that monitors network traffic in real time or reviews the audit logs on servers looking for security violations.

Firewalls filter traffic crossing a border of a network environment in order to block unauthorized or security-violating traffic.

A switch is a networking device used to establish temporary electronically dedicated links between two systems to improve throughput and communications isolation.

A padded cell is a fake network environment generated to lure and snare intruders.

q_ids_ids_03_secp8

An active IDS system often performs which of the following actions? (Select two.)

Answers:

- Requests a second logon test for users performing abnormal activities.
- ***Updates filters to block suspect traffic.**
- ***Performs reverse lookups to identify an intruder.**
- Traps and delays the intruder until the authorities arrive.
- Cannot be detected on the network because it takes no detectable actions.

Explanation:

An active IDS performs behaviors that can be seen by anyone watching the network. Usually, these actions are necessary to block malicious activities or discover the identity of an intruder. Updating filters and performing reverse lookups are common behaviors of an active IDS.

No form of IDS requires users to perform a second logon based on questionable activities. There are some authentication systems, such as CHAP, that periodically re-authenticate, but that is done at random time intervals and is not visible to the user. A solution that serves to trap and delay the intruder until the authorities arrive describes a man trap (a physical security mechanism). However, this definition could be stretched to include honeypots and padded cells (logical or technical security mechanisms often used in conjunction with an IDS).

A passive IDS cannot be detected on the network because it takes no detectable actions.

q_ids_ids_04_secp8

Which of the following devices can monitor a network and detect potential security attacks?

Answers:

- ***IDS**
- Proxy
- CSU/DSU
- DNS server

Explanation:

An intrusion detection system (IDS) is a special network device that can detect attacks and suspicious activity.

A proxy server is a type of firewall that can filter based on upper-layer data.

A CSU/DSU is a device that converts the signal received from the WAN provider into a signal that can be used by equipment at the customer site.

A DNS server provides IP address-to-host name resolution.

q_ids_ids_05_secp8

An IT security manager at a technical college wants to increase the use of controls that generate alerts where ongoing attacks are suspected in the organization's network infrastructure.

Which of the following is a suitable illustration of this type of control?

Answers:

- ***Implementing an intrusion detection system**
- Using strong password policies
- Establishing a secure firewall
- Regularly updating antivirus software

Explanation:

An intrusion detection system (IDS) is a detective control that monitors network traffic and analyzes it for signs of possible intrusions, such as security threats or malicious activities.

Although important, strong password policies are typically preventive controls. They work to prevent unauthorized access by requiring users to set strong passwords but do not actively detect breaches or malicious activities.

Firewalls are primarily preventive controls. They prevent unauthorized access to or from a network by controlling incoming and outgoing traffic based on predetermined security rules. They do not actively detect breaches or malicious activities.

Regularly updating antivirus software is a preventive control. Updated antivirus software helps prevent malware from infecting the system but does not detect breaches or security incidents actively.

q_ids_ids_06_secp8

An e-commerce company recently identified suspicious activity on its web-based application suggesting a zero-day exploit. The security team suspects that a vulnerability in the application might be under active exploitation by malicious actors before the company identified and patched it.

With no known fixes available for a zero-day exploit, what should be the initial course of action for the security team to minimize potential damage and safeguard the application and its users?

Answers:

- ***Implement intrusion detection systems (IDS) and application firewalls.**
- Perform a comprehensive system backup.
- Update the antivirus software on all company devices.
- Enforce password changes for all users.

Explanation:

IDS and application firewalls provide immediate protection by identifying and blocking potentially malicious activity, thus serving as an effective response to zero-day exploits in a web-based application.

Although system backups are important for data recovery, they do not directly address or prevent an active zero-day exploit.

While updating antivirus software can help protect against known threats, it is often ineffective against zero-day exploits as these are previously unknown vulnerabilities that antivirus signatures do not cover yet.

Enforcing password changes for all users might improve account security, but it doesn't directly address the zero-day exploit in the application's code.

q_ids_ips_01_secp8

You are concerned about attacks directed at your network firewall. You want to be able to identify and be notified of any attacks. In addition, you want the system to take immediate action to stop or prevent the attack, if possible.

Which tool should you use?

Answers:

- IDS
- ***IPS**
- Packet sniffer
- Port scanner

Explanation:

Use an intrusion prevention system (IPS) to both detect and respond to attacks.

An intrusion detection system (IDS) can detect attacks and send notifications, but it cannot respond to attacks.

Use a port scanner to check for open ports on a system or a firewall.

Use a packet sniffer to examine packets on the network.

q_ids_ips_02_secp8

Which of the following devices is capable of detecting and responding to security threats?

Answers:

- IDS
- ***IPS**
- DNS server
- Multi-layer switch

Explanation:

An intrusion prevention system (IPS) can detect and respond to security events. An IPS differs from an IDS because it can respond to security threats, not just detect them.

A DNS server provides IP address-to-host name resolution.

A multi-layer switch uses an ASIC module to switch packets based on packet or data content instead of using the CPU and software.

q_ids_ips_03_secp8

Your organization uses a web server to host an e-commerce site.

Because this web server handles financial transactions, you are concerned that it could become a prime target for exploits. You want to implement a network security control that analyzes the contents of each packet going to or from the web server. The security control must be able to identify malicious payloads and block them.

What should you do?

Answers:

- Implement an application-aware IDS in front of the web server.
- ***Implement an application-aware IPS in front of the web server.**
- Install an anti-malware scanner on the web server.
- Implement a packet-filtering firewall in front of the web server.
- Implement a stateful firewall in front of the web server.

Explanation:

You should implement an application-aware IPS in front of the web server. Even though an application-aware IDS can analyze network packets to detect malicious payloads, only an application-aware IPS can both detect and block malicious packets. Because of this, an application-aware IPS would be the most appropriate choice.

Installing an anti-malware scanner on the web server itself is a good idea, but it can only detect malware after it has been installed on the server.

Using a packet-filtering firewall or a stateful firewall is also a good security measure, but neither are capable of inspecting the contents of network packets.

A packet-filtering firewall can only filter based on IP address, port, and protocol.

A stateful firewall can only monitor the state of a TCP connection. These devices should be used in conjunction with an IDS or an IPS to protect a network.

q_ids_logs_anomaly_secp8

A security architect at a multinational corporation designs a comprehensive security strategy to defend against advanced persistent threats (APTs).

The strategy includes real-time analysis of network traffic, detection of unknown threats, correlation of events across different layers, and automatic response to mitigate risks. The architect also ensures compliance with international regulations on data privacy.

Which data sources and security components combination would BEST align with this multifaceted approach?

Answers:

- ***Intrusion prevention system/intrusion detection system (IPS/IDS) logs, anomaly detection, and automated response mechanisms**
- Application logs and endpoint security
- Network logs, firewall logs, and operating system (OS)-specific security logs
- Vulnerability scans, packet captures, and dashboards

Explanation:

Intrusion prevention system/intrusion detection system (IPS/IDS) logs, anomaly detection, and automated response mechanisms provide real-time network traffic analysis, detect unknown threats through anomaly detection, correlate events, and enable automatic response, ideal for advanced security requirements.

Application logs and endpoint security provide insights into software activities and endpoint protection but lack real-time network analysis and advanced threat detection capabilities.

Network, firewall, and operating system (OS)-specific security logs offer information about network operations and connections but don't provide real-time analysis or automatic response to advanced threats.

Vulnerability scans, packet captures, and dashboards provide a view of system weaknesses and network packets but miss real-time threat detection and mitigation capabilities.

q_ids_network_based_secp8

An eCommerce company utilizes an advanced firewall for network protection. However, due to an increasing variety of threats, their existing setup is not enough. The company considers enhancing its firewall with intrusion detection and prevention systems (IDS/IPS).

What should be their primary approach in this scenario?

Answers:

- ***Add network-based IDS/IPS and maintain the firewall rules.**
- Discard firewall rules and rely solely on network-based IDS/IPS.
- Abandon the firewall and adopt host-based IDS/IPS.
- Keep the firewall and ignore IDS/IPS adoption.

Explanation:

Introducing network-based IDS/IPS while keeping firewall rules will augment the company's network protection. The IDS/IPS will monitor for known threats and unusual network packet behavior while the firewall continues filtering traffic based on set rules.

Discarding firewall rules to rely solely on network-based IDS/IPS neglects the benefits of firewall's traffic filtering, resulting in less comprehensive protection.

Abandoning the firewall for host-based IDS/IPS lacks the network-wide threat detection and doesn't leverage the firewall's traffic filtering capabilities.

Keeping the firewall while ignoring IDS/IPS adoption doesn't capitalize on the added security layer that IDS/IPS systems provide. This could result in missed threats.

q_ids_placement_01_secp8

A company has expanded its operations to a new location and is setting up its network infrastructure. A significant part of this setup includes strategically placing devices for optimal security and efficiency.

How should the network security manager decide the optimal placement of the intrusion detection system (IDS) in the new network topology to ensure maximum visibility and efficiency without impacting overall network performance?

Answers:

- Place the IDS outside the firewall.
- Place the IDS at the end of the network.
- ***Place the IDS directly behind the router.**

- Place the IDS near the servers.

Explanation:

Placing the IDS directly behind the router ensures visibility of all incoming and outgoing traffic, which is crucial for detecting any unusual patterns or potential threats.

Placing the IDS outside the firewall exposes it to all internet traffic, including a lot of noise and potential direct attacks on the IDS itself, which is not optimal.

Placing the IDS at the end of the network could result in it only seeing outbound traffic, thus limiting its visibility and effectiveness.

While placing the IDS near the servers could provide some visibility into the traffic going to and from the servers, it will not give it a comprehensive view of all network traffic.

q_ids_placement_02_secp8

A medium-sized organization is upgrading its network infrastructure to secure its enterprise infrastructure by implementing an intrusion prevention system (IPS) and an intrusion detection system (IDS).

The organization has sensitive data in different security zones, and the IT manager has concerns regarding the attack surface and network connectivity.

Which of the following placements of the IPS/IDS devices would be MOST effective in this scenario?

Answers:

- ***Place the IPS/IDS devices at the network perimeter to monitor inbound and outbound traffic.**
- Place the IPS/IDS devices just inside the organization's firewall to monitor the internal network.
- Place the IPS/IDS devices at each end of the VPN tunnel to monitor remote access.
- Place the IPS/IDS devices near the load balancer to monitor traffic distribution.

Explanation:

Placing the IPS/IDS devices at the network perimeter allows for monitoring of all inbound and outbound traffic, providing comprehensive visibility and enabling immediate response to potential threats.

While placing the IPS/IDS devices just inside the firewall can help monitor internal network traffic, it would not provide comprehensive visibility into inbound and outbound traffic.

Although placing the IPS/IDS devices at each end of the VPN tunnel can help monitor remote access, it would not provide complete visibility into all network traffic.

Placing the IPS/IDS devices near the load balancer can help monitor traffic distribution, but it would not provide comprehensive visibility into all network traffic.

q_ids_placement_03_secp8

A large organization is redesigning its network infrastructure to increase security and reduce the potential attack surface. The organization considers implementing an intrusion prevention system (IPS) and an intrusion detection system (IDS) into its security zones.

The IT manager wants to secure connectivity and considers different network appliances and port security measures.

Which of the following options BEST describes the benefits and disadvantages of placing the IPS/IDS devices inline with the network traffic?

Answers:

- ***Inline placement allows for active prevention measures but can become a single point of failure.**
- Inline placement increases the attack surface but provides comprehensive traffic visibility.
- Inline placement reduces connectivity but allows for passive detection.
- Inline placement improves port security but increases network latency.

Explanation:

Inline placement of IPS/IDS devices allows for active prevention measures, such as blocking malicious traffic, but can become a single point of failure if the device malfunctions, potentially disrupting network traffic.

Inline placement does not inherently increase the attack surface, and while it provides comprehensive traffic visibility, this does not address the risk of becoming a single point of failure.

Inline placement does not inherently reduce connectivity, and while it can allow for passive detection, it also allows for active prevention measures not highlighted in this option.

Inline placement does not inherently improve port security, and while it can potentially increase network latency, this does not address the benefit of active prevention measures.

q_ids_placement_04_secp8

A multinational corporation wants to enhance its security infrastructure by deploying an intrusion detection system (IDS) across its global network. The IT manager is considering the placement of IDS sensors to ensure comprehensive network visibility.

Which IDS sensor placement is the MOST effective in this scenario?

Answers:

- ***Place the IDS sensors at network choke points.**
- Place the IDS sensors on the internal network.
- Place the IDS sensors on the external network.
- Place the IDS sensors near the network perimeter.

Explanation:

Placing the IDS sensors at network choke points ensures that they can monitor both inbound and outbound traffic, providing comprehensive visibility across the network. This helps detect malicious activity at the earliest possible stage.

While placing the IDS sensors on the internal network can help monitor outbound traffic, it does not provide comprehensive visibility as it may not capture all inbound traffic.

Putting the IDS sensors on the external network can help monitor inbound traffic but may not provide complete visibility into outbound traffic.

Locating the IDS sensors near the network perimeter can help monitor internal network traffic but may not provide complete visibility into all inbound and outbound traffic.

q_ids_signature_01_secp8

Which IDS method searches for intrusion or attack attempts by recognizing patterns or identifying entities listed in a database?

Answers:

- Heuristics-based IDS
- Anomaly-analysis-based IDS
- Stateful-inspection-based IDS
- ***Signature-based IDS**

Explanation:

A signature-based IDS, or pattern-matching-based IDS, is a detection system that searches for intrusion or attack attempts by recognizing patterns that are listed in a database.

A heuristics-based IDS is able to perform some level of intelligent statistical analysis of traffic to detect attacks. Anomaly-analysis-based IDSs look for changes in the normal patterns of traffic.

Stateful-inspection-based IDSs search for attacks by inspecting packet contents and associating one packet with another.

These searches look for attacks in overall data streams rather than individual packets.

q_ids_signature_02_secp8

What does IDS do when using signature recognition use to identify attacks?

Answers:

- ***Makes comparisons to known attack patterns**
- Alerts for exceeding threshold values
- Statistical analysis to find unusual deviations
- Makes comparison of current statistics to past statistics

Explanation:

Signature recognition, also referred to as pattern matching, dictionary recognition, or misuse-detection (MD-IDS), looks for patterns in network traffic and compares them to known attack patterns called signatures.

Anomaly recognition, also referred to as behavior, heuristic, or statistical recognition, monitors traffic to define a standard activity pattern as normal.

Clipping levels or thresholds are defined that identify deviations from the norm. When the threshold is reached, an alert is generated or an action is taken.

Anomaly-based systems can recognize and respond to some unknown attacks (attacks that do not have a corresponding signature file).

q_ids_signature_03_secp8

You have just installed a new network-based IDS system that uses signature recognition.

What should you do on a regular basis?

Answers:

- ***Update the signature files**
- Check for backdoors
- Generate a new baseline
- Modify clipping levels

Explanation:

Signature recognition, also referred to as pattern matching, dictionary recognition, or misuse-detection (MD-IDS), looks for patterns in network traffic and compares them to known attack patterns called signatures. Signature-based recognition cannot detect unknown attacks; they can only detect attacks identified by published signature files. For this reason, it is important to update signature files on a regular basis.

Anomaly recognition, also referred to as behavior, heuristic, or statistical recognition, monitors traffic to define a standard activity pattern as normal.

Clipping levels, or thresholds, identify deviations from the norm. When the threshold is reached, an alert is generated or an action is taken.

q_ids_signature_04_secp8

A cybersecurity analyst for a large organization is enhancing the company's security posture. The analyst notices increased alerts related to a particular known exploit in the company's server software.

The company's intrusion detection system (IDS) has a rule specifically for this exploit.

In the given scenario, which detection method should the analyst modify or enhance to effectively detect and alert for this specific exploit?

Answers:

- ***Signature-based detection**
- Behavioral-based detection
- Anomaly-based detection
- Trend analysis

Explanation:

Since the exploit is known and the IDS already has a rule set for signature-based detection of this specific exploit, enhancing or focusing on signature-based detection would be the most effective method.

Although behavioral-based detection can detect deviations from normal behavior, it is not the primary method for detecting known threats, especially when a signature for the attack is already in place.

Anomaly-based detection focuses on irregularities in the use of protocols. While useful, it is not the primary method when a known exploit signature is in place.

Trend analysis, while valuable for understanding the environment over time and identifying patterns, does not directly detect specific known threats. Its focus is on long-term patterns rather than immediate threats.

q_ids_signature_anomaly_secp8

How does signature-based detection differ from anomaly-based detection in an intrusion detection system (IDS)?

Answers:

- ***Signature-based detection relies on known attack patterns, while anomaly-based detection identifies deviations from normal behavior.**
- Signature-based detection identifies deviations from normal behavior, while anomaly-based detection relies on known attack patterns.
- Both signature-based and anomaly-based detection rely on known attack patterns.
- Both signature-based and anomaly-based detection identify deviations from normal behavior.

Explanation:

Signature-based detection uses known patterns of malicious behavior (or "signatures") to identify potential threats. This method can quickly and accurately detect known threats but may fail to identify new, unrecorded types of attacks. In contrast, anomaly-based detection works by defining what's considered "normal" behavior within a system and then alerting on any behavior that deviates from this norm.

Signature-based detection relies on known attack patterns, while anomaly-based detection identifies deviations from normal behavior.

Anomaly-based detection identifies threats by noting deviations from established norms of system behavior.

Signature-based detection does not identify deviations from normal behavior. It identifies matches to known attack patterns. Anomaly-based detection identifies normal behavior deviations but does not rely on known attack patterns.

q_ids_threat_secp8

A mid-sized tech company has started experiencing regular system slowdowns and data traffic abnormalities. However, its current intrusion detection system (IDS) has generated no alerts.

The IT department relies heavily on the IDS for potential threats and does not actively monitor system metrics or logs.

Which statement is MOST likely true about the situation?

Answers:

- ***The company is facing a new type of threat not recognized by the IDS.**
- The abnormalities are coincidental and do not signify a potential threat.
- The IDS is fully capable of identifying all potential threats.
- The IDS does not need updating as it handles all kinds of threats.

Explanation:

Despite observable anomalies, the absence of alerts from the IDS suggests that the IDS does not recognize the threat.

While it is possible for system slowdowns or data traffic abnormalities to occur naturally due to factors like system updates or increased workload, the consistent occurrence of these anomalies usually indicates a potential threat, which warrants investigation.

The rapid evolution of cybersecurity threats means that a system can only identify some potential threats, especially if the threats are new or highly sophisticated.

Regular updates are crucial for an IDS, as these updates often contain patches for new threats and vulnerabilities. Without regular updates, an IDS may fail to identify newer threats, leaving the organization vulnerable.

q_ids_trend_analysis_01_secp8

A multinational corporation has recently implemented an intrusion detection system (IDS) and intrusion prevention system (IPS) to protect its network infrastructure.

The security team receives many alerts and struggles to manage false positives. The team must optimize the IDS and IPS to identify and prioritize actual threats while minimizing irrelevant alerts.

Which primary strategy should the team adopt to achieve this objective?

Answers:

- ***Implement trend analysis to identify patterns and anomalies, tune the IDS/IPS over time, and prioritize genuine threats.**
- Ignore all alerts from the IDS/IPS to focus on manual monitoring of network traffic.
- Apply signature-based detection rules only to filter out false positives.
- Integrate SELinux policies for a layered security approach, ensuring system-level restrictions to applications and processes.

Explanation:

Trend analysis is the correct approach, as it allows the security team to understand the environment over time, identify patterns, and distinguish anomalies.

Ignoring all alerts from the IDS/IPS is a poor approach as the company would fail to leverage the benefits of these systems, including its ability to detect and prevent security threats more efficiently than manual monitoring.

Applying only signature-based detection rules would not be sufficient because these rules depend on known patterns of attacks.

Integrating Security Enhanced Linux (SELinux) policies provides an additional layer of protection by enforcing mandatory access controls (MAC) at the operating system level. However, on its own, it doesn't directly optimize the IDS/IPS for reducing false positives.

q_ids_trend_analysis_02_secp8

A global organization recently became the target of a wave of cyberattacks and decided to strengthen its network defenses by deploying an IDS/IPS.

However, the large number of alerts the systems generate cause the security team to spend a significant amount of time dealing with false positives. The team wishes to improve the efficiency of these systems by recognizing and prioritizing actual threats while reducing unnecessary alerts.

Which primary approach should the team employ to meet this goal?

Answers:

- ***Implement trend analysis to identify patterns and anomalies, tune the intrusion detection system/intrusion prevention system over time, and prioritize genuine threats.**
- Ignore all alerts from the intrusion detection system/intrusion prevention system to focus on manual monitoring of network traffic.
- Apply signature-based detection rules only to filter out false positives.
- Integrate SELinux policies for a layered security approach, ensuring system-level restrictions to applications and processes.

Explanation:

Trend analysis is the correct approach, as it allows the security team to understand their environment over time, identify patterns, and distinguish anomalies.

Ignoring all alerts from the intrusion detection system/intrusion prevention system is a poor approach as the company would fail to leverage the benefits of these systems, including its ability to detect and prevent security threats more efficiently than manual monitoring.

Applying only signature-based detection rules would not be sufficient because these rules depend on known patterns of attacks.

Integrating SELinux policies provides an additional layer of protection by enforcing mandatory access controls (MAC) at the operating system level. However, on its own, it doesn't directly optimize the IDS/IPS for reducing false positives.

6.4 Protocol Analyzers

As you study this section, answer the following questions:

- What mode must a NIC be in to perform packet sniffing?
- What needs to be configured on a switch so all packets are sent to the sniffing device?
- Why would a network administrator need to use a protocol analyzer?

In this section, you will learn to:

- Analyze network traffic.

The key terms for this section include:

Term	Definition
Protocol analyzer	Hardware or software used for monitoring and analyzing digital traffic over a network. Protocol analyzers go by other names, such as packet sniffers, packet analyzers, network analyzers, network sniffers, or network scanners.

Promiscuous mode	A mode in which the NIC processes every frame it sees, not just those addressed to it.
Port mirroring	A switch mode in which all frames sent to all other switch ports will be forwarded on the mirrored port.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.4 Given a scenario, analyze indicators of malicious activity.
	• Network attacks
	  ○ On-path
	  ○ Credential replay
	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Hardening techniques
	  ○ Disabling ports/protocols
	4.1 Given a scenario, apply common security techniques to computing resources.
	• Monitoring
	4.5 Given a scenario, modify enterprise capabilities to enhance security.
	• Firewall
	  ○ Ports/protocols
	• Implementation of secure protocols
	4.9 Given a scenario, use data sources to support an investigation.
	• Data sources
	  ○ Packet captures

6.4.1 Protocol Analyzers (Lesson Video)

Transcript:

Protocol analyzers are used to capture data packets moving across a network and analyze them. In most cases, you'll see these tools as software programs that run a computer system, but there are hardware analyzers available as well.

Be aware that protocol analyzers are also called packet analyzers, packet sniffers, network analyzers, network sniffers, or network scanners. These terms all mean roughly the same thing.

Some of the more common protocol analyzers are Wireshark, tcpdump, Cain and Abel, and WinDump. In this lesson, I'll go over how these analyzers work and what we use them for.
The most common way to set up a protocol analyzer is to install the program on a computer that's connected to a network switch. This can be a desktop computer, but laptops are also useful as they provide greater mobility.
When we use a protocol analyzer, the computer's network interface card, or NIC, must be placed in promiscuous mode. Promiscuous mode means that the NIC is able to see all packets on the same network segment. In non-promiscuous mode, the NIC only receives packets that are addressed to its MAC address.
Once we set up and configure our computer with the protocol analyzer, we can begin monitoring the packets that run across the network. We can usually just sit back and monitor the data without much interaction. This is known as passive interception.
Once we actually begin capturing packets, there's quite a few things we can do with this data.
Depending on what we're trying to accomplish, there are different ways we can take advantage of the protocol analyzer's features.
For example, a network administrator may use their protocol analyzer to look for specific protocols like SMTP, DNS, POP3, or ICMP that have packets running across the network. The administrator could even detect an employee's unauthorized internet use by scanning URLs that they find in these packets.
The network SecOps team could use their packet analyzer during a vulnerability assessment. This could reveal open ports or whether passwords are being sent in cleartext or not. They could also detect any malformed or fragmented packets, which would indicate that someone is trying to get around the firewall.
Malicious users can also use a protocol analyzer for many of these same purposes. They could use it to fingerprint a system, which means determining which operating system is running based on how the system responds to different types of network traffic.
Malicious users can also use protocol analyzers to perform active interception. This is when the user intercepts data in order to perform attacks or session hijacking.
MAC flooding is another issue for you to be aware of. A bad actor can connect to a switch and, in short, overload it with fake MAC addresses. MAC flooding essentially turns the switch into a hub and allows the hacker to see all passing traffic.
Most times, though, protocol analyzers are used for valid purposes. Just be aware that unscrupulous people sometimes use them to get around security systems.
That'll wrap up this lesson. In this lesson, we looked at protocol analyzers and how they work. We also went over protocol analyzers' many uses and misuses, and we saw how you can utilize this tool for your role in system security. Always remember to put the NIC in promiscuous mode to properly capture data when choosing to use a protocol analyzer.

6.4.2 Protocol Analyzer Facts

A protocol analyzer is hardware or software for monitoring and analyzing digital traffic over a network. Protocol analyzers go by other names, such as packet sniffers, packet analyzers, network analyzers, network sniffers, or network scanners.

This lesson covers the following topics:

- How protocol analyzers work
- Uses of protocol analyzers

How Protocol Analyzers Work

Typically, a protocol analyzer is run on a single device with the intent of capturing frames for all other devices on the network or subnet. Using a packet sniffer this way requires the following configuration changes:

- A network interface card, or NIC, accepts frames addressed only to that NIC by default. To enable the packet sniffer to capture frames sent to other devices, configure the NIC in promiscuous mode (sometimes called p-mode). In p-mode, the NIC processes every frame it sees.
- When using a switch, the switch only forwards packets to the switch port that holds a destination device. This means that when your packet sniffer is connected to a switch port, it does not see traffic sent to other

switch ports. To configure the switch to send all frames to the packet sniffing device, configure port mirroring on the switch. With port mirroring, all frames sent to all other switch ports are forwarded on the mirrored port.

A malicious user that is connected to the switch can accomplish the same thing by running a MAC flooding attack. This attack essentially turns the switch into a hub, and traffic is sent to all ports.

By default, a protocol analyzer is a passive device. It copies frames and allows viewing of frame contents, but it does not allow the capture, modification, or re-transmission of frames. This is referred to as passive interception.

When using a protocol analyzer, the user can filter the frames so they see only the frames with information of interest. You can save the results of a capture to analyze frames at a later time or on a different device.

- Filters show only those frames or packets to or from specific addresses or frames that include specific protocol types.
- A capture filter captures (records) only the frames identified by the filter. Frames not matching the filter criteria are not captured.
- A display filter shows only the frames that match the filter criteria. Frames not matching the filter criteria are still captured but not shown.

Uses of Protocol Analyzers

Depending on the role of the user, a protocol analyzer may have different uses. The following table shows some common roles and how they may use a protocol analyzer:

Role	Protocol Analyzer Usage

Network administrator	A network administrator can use the protocol analyzer to assist in the management of the network and employee usage. The protocol analyzer can help to:
	• Monitor and log network traffic as it is transmitted over the network.
	• Check for specific protocols on the network, such as SMTP, DNS, POP3, and ICMP. Identifying the specific protocols helps to:
	  ○ Identify devices that might be using unallowed protocols, such as ICMP, or legacy protocols, such as IPX/SPX or NetBIOS.
	  ○ Identify traffic that might be sent by attackers.
	• Examine the data contained within a packet. For example, by looking at the packet data, the network administrator can identify users connecting to unauthorized websites.
	• Analyze network performance
	• Troubleshoot communication problems or investigate the source of heavy network traffic

Security operations	The network SecOps team can use the protocol analyzer during a vulnerability assessment. The protocol analyzer can help the SecOps team to: • Identify frames that might cause errors. For example, the network administrator can: ○ Determine which flags are set in a TCP handshake

	○ Detect any malformed or fragmented packets. This would indicate that someone is trying to get around the firewall. • Discover passwords and other sensitive data being sent in cleartext. • Find any open network ports that should not be open.
Malicious user/hacker	A malicious user can use the protocol analyzer to find the same information as the network administrator and SecOps teams. By themselves, protocol analyzers cannot be used to perform an attack. However, protocol tools can be used with protocol analyzers for active interception of network traffic to perform attacks, such as: • Spoofing • Man-in-the-middle attacks • Replay attacks • TCP/IP session hijacking • MAC flooding A hacker can also use the analyzer to perform system fingerprinting. System fingerprinting identifies which operating system the system is running based on how it responds to different types of network traffic.

Common protocol analyzers include:

- Wireshark
- Ethereal
- Dsniff
- Ettercap
- Tcpdump
- Windump
- Cain and Abel

6.4.3 Analyzing Network Traffic (Demo Video)

Transcript:

Wireshark is a network packet analyzer that tries to capture network packets and display the data they carry in as much detail as possible.

Network professionals use Wireshark to troubleshoot network problems, examine security problems, verify network devices, debug network issues, and more. Wireshark can be installed on Windows, Linux, and other operating systems. It's not a firewall or intrusion detection system, and it doesn't keep bad things from happening. But it does a great job of monitoring and measuring activity on a network.
We're going to take a very brief look at Wireshark. There are entire courses on Wireshark, and as you progress professionally, you'll want to learn more about the program.
Let's start with Wireshark's user interface. Throughout this course we'll use Wireshark on Linux and Windows. I'm currently on Kali Linux machine. We're going to use Wireshark from there. Like most programs, Wireshark has a menu at the top. The File menu is similar to other programs, as you can see when we click on it. The same is true for the other menu items across the top.
The main toolbar gives you quick access to the most common Wireshark tasks. The two most common items here are the shark fin, or Start, which starts capturing packets, and this square, which turns red when Wireshark is capturing.

Below that, there's the Filter toolbar. We'll come back to this in a few minutes.
Down here, there's the Packet List pane. This is where all the captured packets are displayed. It only takes a few minutes to see hundreds or thousands of packets, depending on your network traffic. Each line here is a packet that Wireshark has captured, and it's separated in columns up here. You can separate packets by source or destination IP addresses or categorize them in other ways. Whichever one is selected up here will display more detailed information down in the Packet Detail pane.
The Packet Detail pane shows all the details about the selected packet. I can come down here, expand the different areas, and see more details. For those of you who love the OSI model, you'll be very happy to hear that this is listed in the OSI order.
Farther down, we have the Packet Bytes pane. This is a hexadecimal dump of the selected packets. Sometimes, you can see readable information over here, in ASCII. If the bytes aren't readable, this information is replaced with periods.
And finally, at the bottom, we have the status bar.
A lot of the Wireshark interface can be customized, but this is what it looks like with the default settings.
Okay. Let's go back to display filters. When you do some sniffing with Wireshark, you're going to be overwhelmed with all of the packets it will capture. We can filter the ones that aren't needed. I already ran a scan earlier, and my traffic is here.
One way to filter is to use the Expression Builder. For example, let's say I want to filter dhcp. If I look, bootp and dhcp are actually the same thing. If I expand this out, I get lots of things to choose from. I'm going to just leave it as-is.
Now let's go over to the Relation box. We have things such as is present. The double equal sign (==) means equal to. The exclamation and equal (!=) means not equal to. The next several should be self-explanatory, but I want to pick the equal to and then, for the value, type in '10.10.10.1', which is my DHCP server. Now, I'll click OK and see what happens. I'm filtering only my DHCP traffic.
Now, I want to see if there are any cleartext usernames and passwords here. I can do that several ways. I can create a filter by typing in 'http contains admin'. As I type, notice the salmon color. As soon as it's a valid filter, it turns green.
Let's come down here and choose this packet. I'm picking it because it has an external public IP address. I'm going to right-click and select Follow TCP Stream. Now it displays the entire stream related to this conversation. Some of it is readable; some is not. I'm looking for the word Admin. Let's see if we can just do a find for that. I'll click on Find Next until I see something worth looking at. Here, it looks like I found something. It's the word log with the username Admin listed. This is followed by pwd and the password the person typed in. So, with Wireshark, we were able to see a cleartext password. We could also filter other things, like FTP, to find cleartext passwords.
I'll come up and clear this filter.
Another way we can filter is to come down here. If I find something that I think is interesting, I can right-click and apply a filter.
We could spend an hour looking at filters, but let's keep moving on here.
So, we kind of did this already. We know we can look at packets. But what if we wanted to look at the entire conversation between two devices? I can come down here, right-click on one of these, and select Follow TCP Stream. Now, I get this window that pops up with the entire stream. This particular stream isn't very useful, but you get the idea.
So, how many different conversations have been going on here? I can figure that out by coming up to Statistics and selecting Conversations. When I do, my conversations are combined. All the packets are organized neatly. I can choose a few different ways to see my data and then sort by the total size of the bytes, largest to smallest. I have some tabs up here, such as IP versions 4 and 6. I can select TCP to display my TCP conversation. Okay. I'll close this and then clear my filter.
So far, when we look at our packets from Wireshark, we see a bunch of gobbledygook. The good news is that Wireshark has a way to export objects. In our case, we want a file. To do this, we can come up here, to File, go down to Export Objects, and select a category. I can pick HTTP, and I'll get a list of things that downloaded from the web, or we can get objects from our local Windows network by selecting SMB. That's the one I'm going to choose. I only have one file. It's one that I transferred across the network earlier. It's a jpg image file. I'll click on Save and pick a place to save it. I'll pick Desktop, click Save, and close this window.
Now I'll minimize Wireshark, go over to my folders, go to Desktop, and here's the picture that I captured in Wireshark.
That's it for this demo. In this demo, we took a look at Wireshark. We started by going over the Wireshark interface. Then we talked about using filters. We illustrated how to follow conversations in Wireshark. Finally, we ended by exporting an object.

6.4.4 Practice Questions (Section Quiz)

q_prot_analyzers_filter_secp8

You are using a protocol analyzer to capture network traffic. You want to only capture the frames coming from a specific IP address.

Which of the following can you use to simplify this process?

Answers:

- ***Capture filters**
- Switch
- NIC
- Display filters

Explanation:

A capture filter records only the frames identified by the filter. Frames that don't match the filter criteria are not captured.

A switch connects multiple computers together in a network. It is not used to capture specific frames.

A network interface card (NIC) is used to transmit and receive frames addressed to it. It is not used to capture specific frames.

A display filter shows only the frames that match the filter criteria. Frames that don't match the filter criteria are still captured, but are not shown.

q_prot_analyzers_finger_secp8

Which of the following processes identifies an operating system based on its response to different types of network traffic?

Answers:

- Port scanning
- ***Fingerprinting**
- Firewalking
- Social engineering

Explanation:

A hacker can use an analyzer to perform system fingerprinting. System fingerprinting identifies which operating system the system is running based on how it responds to different types of network traffic.

Port scanning pings every port on an external interface or attempts a connection in order to discover which ports are open and active, and which ones are not.

Firewalking uses the traceroute command to discover which services can pass through a firewall or router.

Social engineering exploits human nature to obtain information. A hacker often impersonates someone of authority and requests data.

q_prot_analyzers_mirroring_secp8

You decide to use a packet sniffer to identify the type of traffic sent to a router. You run the packet sniffing software on a device that is connected to a hub with three other computers. The hub is connected to a switch that is connected to the router.

When you run the software, you see frames addressed to the four workstations, but not to the router.

Which feature should you configure on the switch?

Answers:

- Promiscuous mode
- ***Port mirroring**
- Bonding
- Spanning Tree Protocol

Explanation:

A switch only forwards packets to the switch port that holds a destination device. This means that when your packet sniffer is connected to a switch port, it does not see traffic sent to other switch ports. To configure the switch to send all frames to the packet sniffing device, configure port mirroring on the switch. With port mirroring, all frames sent to all other switch ports are forwarded on the mirrored port.

Promiscuous mode configures a network adapter to process every frame it sees, not just the frames addressed to that network adapter. In this scenario, you know that the packet sniffer is running in promiscuous mode because it can already see frames sent to other devices.

Bonding logically groups two or more network adapters together to be used at the same time for a single logical network connection.

Spanning Tree Protocol (STP) runs on a switch and ensures that there is only one active path between switches, allowing for backup-redundant paths.

q_prot_analyzers_output_01_secp8

You are running a packet sniffer on your workstation so you can identify the types of traffic on your network. You expect to see all the traffic on the network, but the packet sniffer only seems to be capturing frames that are addressed to the network interface on your workstation.

Which of the following must you configure in order to see all of the network traffic?

Answers:

- ***Configure the network interface to use promiscuous mode.**
- Configure the network interface to use port mirroring mode.
- Configure the network interface to use protocol analysis mode.
- Configure the network interface to enable logging.

Explanation:

Configure the network interface to use promiscuous mode. By default, a NIC only accepts frames addressed to itself. To enable the packet sniffer to capture frames sent to other devices, configure the NIC in promiscuous mode (sometimes called p-mode). In p-mode, the NIC processes every frame it sees.

Configuring port mirroring on a switch sends all frames to the packet sniffing device. With port mirroring, all frames sent to all other switch ports are forwarded on the mirrored port. Configuring port mirroring would not resolve the issue limited capture of frames.

Network protocol analysis is a network sniffer that captures data for further analysis and understanding. Configuring protocol analysis would only capture and save frames. It would not resolve the issue of limited capturing of frames.

Configuring the network interface to enable logging would provide an additional block to capturing frames that are addressed to the network interface on your workstation.

q_prot_analyzers_output_02_secp8

Which of the following accurately describes what a protocol analyzer is used for? (Select two.)

Answers:

- ***A passive device that is used to copy frames and allow you to view frame contents.**
- ***A device that does NOT allow you to capture, modify, and retransmit frames (to perform an attack).**
- A device that allows you to capture, modify, and retransmit frames (to perform an attack).
- A device that can simulate a large number of client connections to a website, test file downloads for an FTP site, or simulate large volumes of emails.
- A device that measures the amount of data that can be transferred through a network or processed by a device.

Explanation:

A protocol analyzer is a passive device that copies frames and allows you to view frame contents. However, it does not allow you to capture, modify, and retransmit frames (activities that are used to perform an attack).

A load tester simulates a load on a server or service. For example, a load tester might simulate a large number of client connections to a website, test file downloads for an FTP site, or simulate large volumes of emails.

A throughput tester measures the amount of data that can be transferred through a network or processed by a device.

q_prot_analyzers_output_03_secp8

You want to identify traffic that is generated and sent through a network by a specific application running on a device.

Which tool should you use?

Answers:

- ***Protocol analyzer**
- Certifier
- Toner probe
- TDR
- Multimeter

Explanation:

Use a protocol analyzer (also called a packet sniffer) to examine network traffic. You can capture or filter packets from a specific device or packets that use a specific protocol.

Use a time-domain reflector (TDR) to measure the length of a cable or to identify the location of a fault in the cable.

A toner probe is two devices used together to trace the end of a wire from a known endpoint into the termination point in the wiring closet.

A cable certifier is a multi-function tool that verifies that a cable or an installation meets the requirements for a specific architectural implementation.

A multimeter is a device that tests various electrical properties, such as voltage, amps, and ohms.

q_prot_analyzers_sniffer_01_secp8

You want to know which protocols are being used on your network. You'd like to monitor network traffic and sort traffic by protocol.

Which tool should you use?

Answers:

- Port scanner
- ***Packet sniffer**
- IDS
- IPS
- Throughput tester

Explanation:

A packet sniffer is special software that captures (records) frames that are transmitted on a network. Use a packet sniffer to:

- Identify the types of traffic on a network.
- View the exchange of packets between communicating devices. For example, you can capture frames related to the domain name system (DNS) and view the exact exchange of packets for a specific name resolution request.
- Analyze packets sent to and from a specific device.
- View packet contents.

Use a port scanner to identify protocol ports that are opened in a firewall or active on a device. A port scanner checks individual systems, while a packet sniffer watches traffic on the network.

A throughput tester measures the amount of data that can be transferred through a network or processed by a device (such as the amount of data that can be retrieved from a disk in a specific period of time).

An IDS is a special network device that can detect attacks and suspicious activity. A passive IDS monitors, logs, and detects security breaches, but it takes no action to stop or prevent the attack. An active IDS (also called an intrusion protection system, or IPS) performs the functions of an IDS, but it can also react when security breaches occur.

q_prot_analyzers_sniffer_02_secp8

You are concerned about attacks directed against the firewall on your network. You would like to examine the content of individual frames sent to the firewall.

Which tool should you use?

Answers:

- ***Packet sniffer**
- Load tester
- Throughput tester
- Event log
- System log

Explanation:

A packet sniffer is special software that captures frames transmitted on the network. Use a packet sniffer to:

- View packet contents.
- Identify the types of traffic on a network.
- View the exchange of packets between communicating devices. For example, you can capture frames related to the domain name system and view the exact exchange of packets for a specific name resolution request.
- Analyze packets sent to and from a specific device.

A load tester simulates a load on a server or service.

A throughput tester measures the amount of data that can be transferred through a network or processed by a device (such as the amount of data that can be retrieved from a disk in a specific period of time).

System and event logs record what has happened on a device. They do not record individual frames or packets.

q_prot_analyzers_soar_secp8

Which of the following roles would be MOST likely to use a protocol analyzer to identify frames that might cause errors?

Answers:

- Network administrator
- Malicious hacker
- Standard user
- ***Security operations team**

Explanation:

The network security operations (SecOps) team can use a protocol analyzer during a vulnerability assessment. The protocol analyzer can help the SecOps team to:

- Identify frames that might cause errors. For example, the network administrator can:
  - Determine which flags are set in a TCP handshake.
  - Detect any malformed or fragmented packets. This would indicate that someone is trying to get around the firewall.
- Discover passwords and other sensitive data being sent in cleartext.
- Find any open network ports that should not be open.

A network administrator can use a protocol analyzer to assist in the management of the network and employee usage. However, a network administrator would not be the most likely to use a protocol analyzer to identify frames that might cause errors.

A malicious hacker could use a protocol analyzer to identify frames that might cause errors, but they most likely would not use it for that purpose.

A standard user should not be using a protocol analyzer on a network for any reason.

q_prot_analyzers_wireshark_secp8

You want to use a tool to see packets on a network, including the source and destination of each packet.

Which tool should you use?

Answers:

- ***Wireshark**
- OVAL
- Nessus
- nmap

Explanation:

A protocol analyzer, also called a packet sniffer, is special software that captures (records) frames that are transmitted on a network. A protocol analyzer is a passive device. It copies frames and allows you to view frame contents, but it does not allow you to capture, modify, and retransmit frames (activities that are used to perform an attack). Wireshark is a popular protocol analyzer.

The nmap command is a tool that performs ping scans (finding devices on the network) as well as port scans (looking for open ports on the network).

Nessus is a vulnerability-scanning tool. While a protocol analyzer looks at packets on the network, a vulnerability scanner looks for weaknesses in systems, including open ports, running services, and missing patches.

Open Vulnerability and Assessment Language (OVAL) is an international standard for testing, analyzing, and reporting a system's security vulnerabilities.

6.5 Analyzing Network Attacks

As you study this section, answer the following questions:

- In which type of attack does the hacker place themselves between two devices to intercept communications?
- What are common indicators of a DDoS attack?
- What is the usual result of a distributed denial-of-service attack?
- What is a reflected DDoS attack?
- What are some indicators of a DNS attack?

In this section, you will learn to:

- Analyze ARP poisoning.
- Poison ARP and analyze with Wireshark.
- Analyze DNS poisoning.
- Poison DNS.
- Analyze SYN flood.

- Perform and analyze a SYN flood.
- Examine DNS attacks.

The key terms for this section include:

Term	Definition
On-path attack	An attack where the threat actor makes an independent connection between two victims and is able to read and possibly modify traffic.
Credential replay	An attack that uses a captured authentication token to start an unauthorized session without having to discover the plaintext password for an account.
Distributed reflected DoS (DRDoS)	A malicious request to a legitimate server is created and sent as a link to the victim, so that a server-side flaw causes the malicious component to run on the target's browser.
Amplification attack	A type of reflected attack that targets weaknesses in specific application protocols to make the attack more effective at consuming target bandwidth. Amplification attacks exploit protocols that allow the attacker to manipulate the request in such a way that the target is forced to respond with a large amount of data.
DNS poisoning	An attack where a threat actor injects false resource records into a client or server cache to redirect a domain name to an IP address of the attacker's choosing.
Distributed denial-of-service (DDoS)	An attack that involves the use of infected Internet-connected computers and devices to disrupt the normal flow of traffic of a server or service by overwhelming the target with traffic.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.1 Compare and contrast common threat actors and motivations.
	2.1.3 - Motivations
	2.1.3.1 - Data exfiltration
	
	2.4 Given a scenario, analyze indicators of malicious activity.
	
	• Network attacks
	  ○ Distributed denial-of-service (DDoS)
	  ▪ Amplified
	  ▪ Reflected
	  ○ Domain Name System (DNS) attacks
	  ○ Wireless
	  ○ On-path
	  ○ Credential replay

	○ Malicious code • Application attacks ○ Privilege escalation 4.7 Explain the importance of automation and orchestration related to secure operations. • Use cases of automation and scripting 4.9 Given a scenario, use data sources to support an investigation. • Data sources ○ Packet captures 5.5 Explain types and purposes of audits and assessments. • Penetration testing ○ Reconnaissance
TestOut Security Pro	5.2 Assessment Techniques • Analyze Network Attacks

6.5.1 Analyzing Network Attacks (Lesson Video)

Transcript:

In this lesson, we'll focus on various network attack techniques and indicators. A network attack is a general category for a number of strategies and techniques that threat actors use to either disrupt or gain access to systems via a network vector. Let's dive in.

We'll begin by discussing distributed denial-of-service, or DDoS, Attacks. A denial-of-service (DoS) attack is anything that reduces the availability of a resource.
It can target physical hardware, like servers and routers, or exploit software vulnerabilities to disrupt services.
Malware-based DoS attacks can even destroy file systems or overload hardware components like CPU, memory, storage, or network bandwidth.
DDoS attacks are a specific breed of DoS attacks. They're called "distributed" because they're launched from multiple sources simultaneously. Threat actors compromise machines and use them as handlers in a command-and-control network, forming a botnet. These botnets are then used to carry out DDoS attacks.
Assembling and managing a botnet large enough to overwhelm a network that has effective DDoS mitigation measures can be a costly endeavor. This has prompted threat actors to devise DDoS techniques that increase the effectiveness of each attack.
In a distributed reflected DoS, or DRDoS, attack, the threat actor spoofs the victim's IP address and attempts to open connections with multiple third-party servers. Those servers direct their SYN/ACK responses to the victim host. This rapidly consumes the victim's available bandwidth.
An amplification attack is a type of reflected attack that targets weaknesses in specific application protocols to make the attack more effective at consuming target bandwidth. Amplification attacks exploit protocols that allow the attacker to manipulate the request in such a way that the target is forced to respond with a large amount of data. Recognizing DDoS attacks is crucial, but it can be challenging.

Look for unexplained traffic spikes with no legitimate explanation. To mitigate these attacks, we often need high-availability services like load balancing and cluster services.
However, DDoS attacks often involve spoofed source addresses, making it challenging to stop them at the source.
Now, let's shift gears to on-path Attacks.
In on-path attacks, threat actors position themselves between two hosts, capturing, monitoring, and even modifying their communication. This means the hosts might not even realize they're being intercepted.
Another attack focuses on the domain name system. The domain name system, or DNS, resolves requests for named hosts and services to IP addresses. Name resolution is a critical addressing method on the internet and on private networks. There are many potential attacks against DNS.
On the public internet, attackers might use typosquatting to confuse victims with malicious sites.
DNS can also be exploited in DRDoS attacks or by hijacking public DNS servers.
If threat actors have access to the same network as the victim, they can use DNS-based on-path attacks. They respond to DNS queries with spoofed replies, often in conjunction with denial-of-service attacks on the victim's legitimate DNS server. In some cases, rogue DHCP servers might be used to redirect clients to DNS resolvers controlled by attackers.
Another DNS attack vector is DNS client cache poisoning. Here, attackers manipulate the HOSTS file, which is a text file containing a known name: IP address mapping. This cache is checked before DNS resolution is used. If this cache is poisoned, it can redirect traffic wherever the attacker chooses. Think of it as altering road signs to lead people astray.
DNS server cache poisoning aims at corrupting the DNS server's records. Attackers can perform DoS against the server holding authorized records and spoof replies to other name servers. DNS server cache poisoning can also involve making the victim DNS server respond to recursive queries with false records.
DNS can also be misused for command and control of remote access Trojans.
To detect DNS attacks, monitoring DNS event logs is crucial. Look out for unusual query types, suspicious IP address ranges, or spikes in DNS lookup failures.
Wireless attacks are security threats targeting wireless networks. Let's look at a few. A rogue access point is an unauthorized network access point connected to a network. It can be set up maliciously or accidentally. An evil twin is a rogue access point mimicking a legitimate one, often using similar SSID names to deceive users.
For example, a malicious user sets up an evil twin with the same SSID as a coffee shop's Wi-Fi, intercepting users' data when they unknowingly connect to it.
A wireless denial-of-service attack prevents clients from connecting to the legitimate access point.
This attack disrupts wireless networks by interfering with signals or targeting clients. Attackers can jam legitimate networks with rogue access points or send spoofed frames to disconnect clients.
The final network attack technique we'll look at in this lesson is a wireless replay attack. A replay attack is an attack where the threat actor intercepts some authentication data and reuses it to try to reestablish a session. A credential replay is an attack that uses a captured authentication token to start an unauthorized session without having to discover the plaintext password for an account.
Well, that's it for this lesson. In this lesson, we went over several network attack strategies. First, we looked at various denial-of-service attacks. We then reviewed on-path attacks. Next, we discussed DNS attacks. And we finished the lesson by looking at wireless attacks. Being familiar with the variety of ways networks can be attacked is one of the most critical aspects of cybersecurity.

6.5.2 Analyzing Network Attacks Facts

This lesson covers the following topics:

- Network attacks
- Distributed denial-of-service (DDoS) attacks
- On-path attacks
- Domain Name System (DNS) attacks

Network Attacks

A network attack is a general category for a number of strategies and techniques that threat actors use to either disrupt or gain access to systems via a network vector. Network attack analysis is usually informed by considering the place each attack type might have within an overall cyberattack lifecycle:

- Reconnaissance is where a threat actor uses scanning tools to learn about the network. Host discovery identifies which IP addresses are in use. Service discovery identifies which TCP or UDP ports are open on a given host. Fingerprinting identifies the application types and versions of the software operating each port and potentially the operating system running on the host and its device type. Rapid scanning generates a large amount of distinctive network traffic that can be detected and reported as an intrusion event. Still, it is very difficult to differentiate malicious scanning activity from non-malicious scanning activity.
- Credential harvesting is a type of reconnaissance where the threat actor attempts to learn passwords or cryptographic secrets that will allow them to obtain authenticated access to network systems.
- Denial-of-service (DoS) in a network context refers to attacks that cause hosts and services to become unavailable. This type of attack can be detected by monitoring tools that report when a host or service is not responding or is suffering from abnormally high volumes of requests. A DoS attack might be launched as an end in itself or to facilitate the success of other types of attacks.
- Weaponization, delivery, and breach refer to techniques that allow a threat actor to get access without having to authenticate. This typically involves various types of malicious code being directed at a vulnerable application host or service over the network or sending code concealed in file attachments and tricking a user into running it.
- Command and control (C2 or C&C), beaconing, and persistence refer to techniques and malicious code that allow a threat actor to operate a compromised host remotely and maintain access to it over a period of time. The threat actor has to disguise the incoming command and outgoing beaconing activity as part of the network's regular traffic, such as by using encrypted HTTPS connections. Detection of this type of activity usually depends on identifying anomalous connection endpoints, such as connections to IP addresses in countries that do not respect copyright or privacy laws. There can also be indicators on the compromised host, such as the malware itself and unauthorized startup items.
- Lateral movement, pivoting, and privilege escalation refer to techniques that allow the threat actor to move from host to host within a network or from one network segment to another and to obtain wider and higher permissions for systems and services across the network. These types of attacks are detected via anomalous account logins and privilege use. Still, detection usually depends on machine learning-backed software, as it is typically difficult to differentiate anomalous behavior from normal behavior.
- Data exfiltration refers to obtaining an information asset and copying it to the attacker's remote machine. Anomalous large data transfers might be an indicator of exfiltration, but a threat actor could perform the attack stealthily by only moving small amounts of data at any one time.

Note that stages in the lifecycle are iterative. For example, a threat actor might perform external reconnaissance and credential harvesting or breach to obtain an initial foothold. They might then perform reconnaissance and credential harvesting from the foothold to perform lateral movement and privilege escalation on internal hosts.

Distributed Denial-Of-Service (DDoS) Attacks

A denial-of-service (DoS) attack is anything that reduces the availability of a resource. DoS attacks can target physical hardware and infrastructure. A malware-based DoS attack might destroy a file system or engineer excessive CPU, memory, storage, or network bandwidth consumption.

A DoS attack can also exploit protocol or configuration weaknesses at different network layers. DoS attacks against network hosts and gateways are typically of a type called distributed DoS (DDoS). DDoS means that the attack is launched from multiple hosts simultaneously. Typically, a threat actor will compromise machines to use as handlers in a command and control network. The handlers are used to compromise thousands or millions of hosts with DDoS bot tools, forming a botnet.

Some types of DDoS attacks simply aim to consume network bandwidth, denying it to legitimate hosts by using overwhelming numbers of bots making ordinary requests. Others cause resource exhaustion on the victim host by bombarding them with requests, which consume CPU cycles and memory. This delays the processing of legitimate traffic and could potentially crash the host system completely. For example, a SYN flood attack works by withholding the client's ACK packet during TCP's three-way handshake. A server, router, or firewall can maintain a queue of pending connections recorded in its state table. When it does not receive an ACK packet from the client, it resends the SYN/ACK packet a set number of times before timing out the connection. The problem is that a server may only be able to manage a limited number of pending connections, which the DDoS attack quickly fills up. This means that the server is unable to respond to genuine traffic.

Reflected Attacks

Assembling and managing a botnet large enough to overwhelm a network that has effective DDoS mitigation measures can be a costly endeavor. This has prompted threat actors to devise DDoS techniques that increase the effectiveness of each attack. In a distributed reflected DoS (DRDoS) attack, the threat actor spoofs the victim's IP address and attempts to open connections with multiple third-party servers. Those servers direct their SYN/ACK responses to the victim host. This rapidly consumes the victim's available bandwidth.

An asymmetric threat is one where the threat actor is able to perpetrate effective attacks despite having fewer resources than the victim.

Amplified Attacks

An amplification attack is a type of reflected attack that targets weaknesses in specific application protocols to make the attack more effective at consuming target bandwidth. Amplification attacks exploit protocols that allow the attacker to manipulate the request in such a way that the target is forced to respond with a large amount of data. Protocols commonly targeted include Domain Name System (DNS), N etwork T ime P rotocol (NTP), and C onnectionless L ightweight D irectory A ccess P rotocol (CLDAP). Another example of a particularly effective attack exploits the Memcached database caching system used by web servers.

DDoS Indicators

DDoS attacks can be diagnosed by traffic spikes that have no legitimate explanation. Still, they can usually only be mitigated by providing high-availability services, such as load balancing and cluster services. In some cases, a stateful firewall can detect a DDoS attack and automatically block the source. However, for many of the techniques used in DDoS attacks, the source addresses will be randomly spoofed or launched by bots, making it difficult to stop the attack at the source.

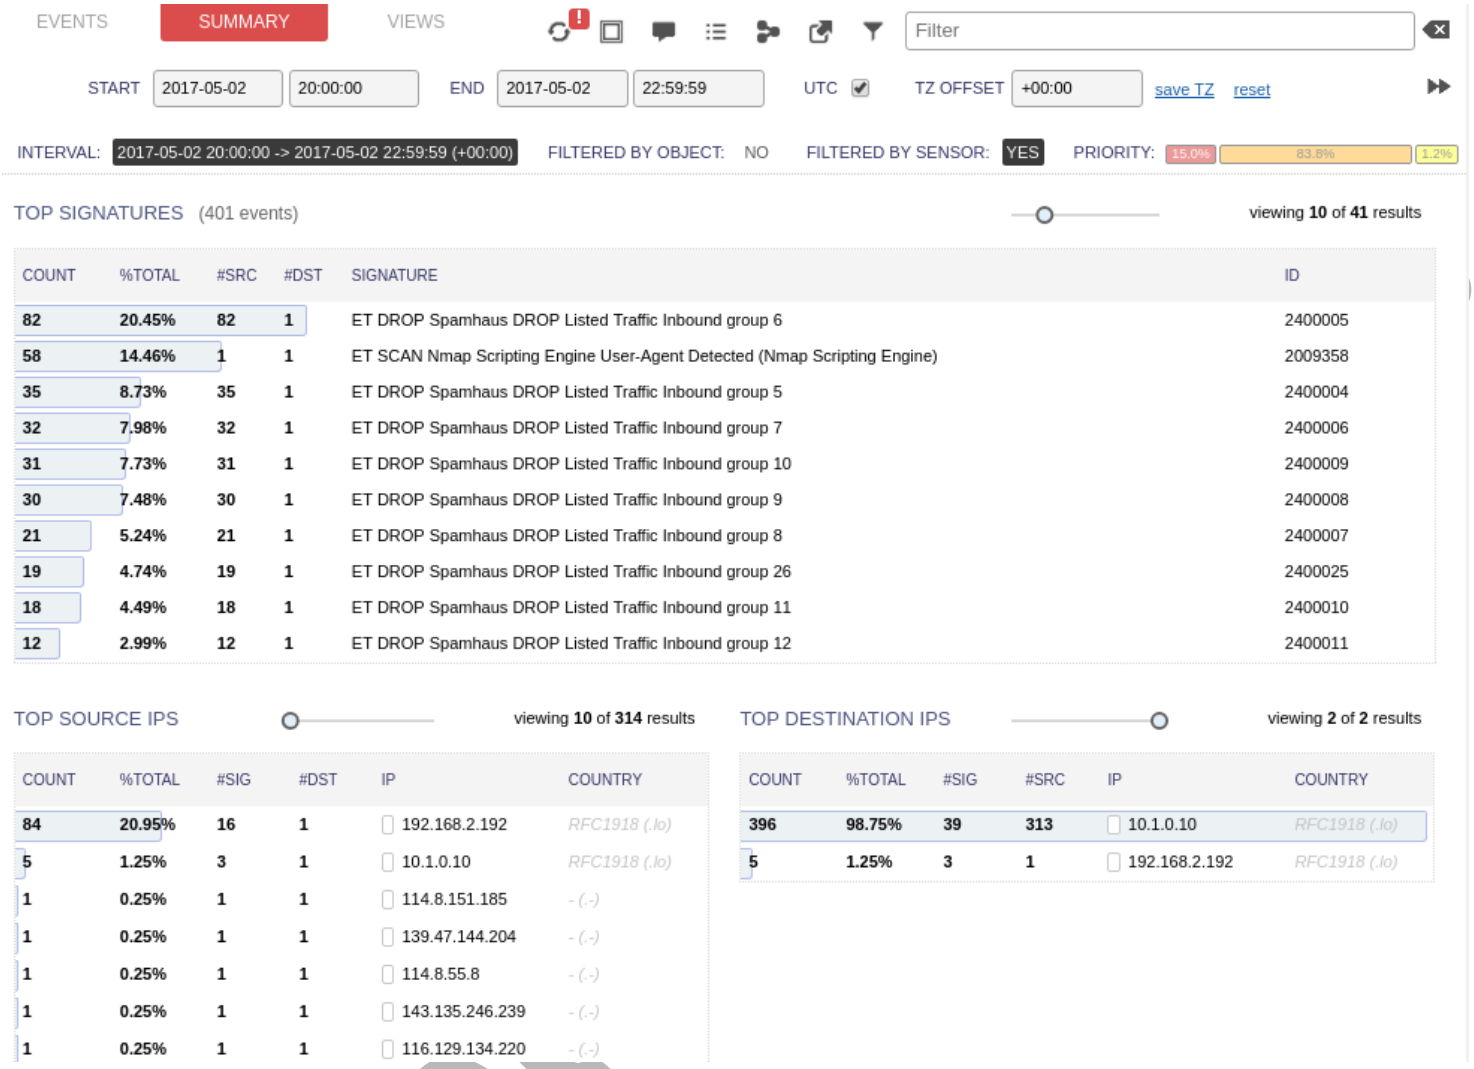

Dropping traffic from blocklisted IP ranges using Security Onion IDS. (Screenshot used with permission from Security Onion.)

On-Path Attacks

An on-path attack is where the threat actor gains a position between two hosts and transparently captures, monitors, and relays all communication between them. Because the threat actor relays the intercepted communications, the hosts might not be able to detect the presence of the threat actor. An on-path attack could also be used to covertly modify the traffic. For example, an on-path host could present a workstation with a spoofed website form to try to capture the user credential. This attack is also referred to as an on-path attack or as an adversary-in-the-middle (AitM) attack.

On-path attacks can be launched at any network layer. One infamous example attacks the way layer 2 forwarding works on local segments. The Address Resolution Protocol (ARP) identifies the MAC address of a host on the local segment that owns an IPv4 address. An ARP poisoning attack uses a packet crafter, such as Ettercap, to broadcast unsolicited ARP reply packets. Because ARP has no security mechanism, the receiving devices trust this communication and update their MAC:IP address cache table with the spoofed address.

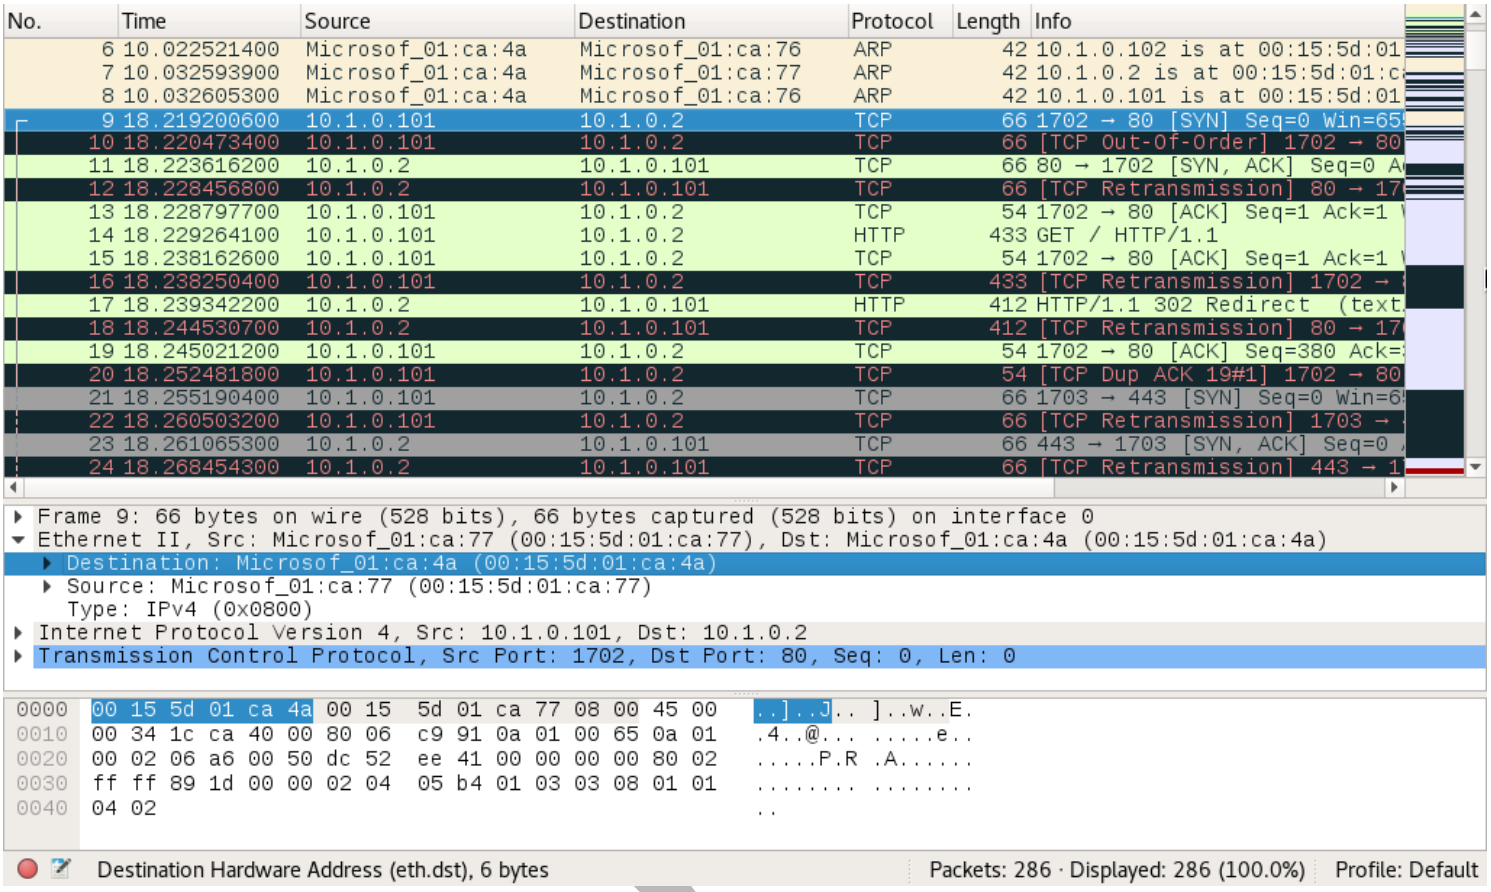

Packet capture opened in Wireshark showing ARP poisoning. (Screenshot used with permission from wireshark.org.)

This screenshot shows packets captured during a typical ARP poisoning attack:

- In frames 6–8, the attacking machine (with MAC address ending :4a) directs gratuitous ARP replies at other hosts (:76 and :77), claiming to have the IP addresses .2 and .102. This pattern of gratuitous ARP traffic is an indicator of the attack.
- In frame 9, the .101/:77 host tries to send a packet to the .2 host, but it is received by the attacking host (with the destination MAC :4a).
- In frame 10, the attacking host re-transmits frame 9 to the actual .2 host. Wireshark colors the frame black and red to highlight the re-transmission.
- In frames 11 and 12, you can see the reply from .2, received by the attacking host in frame 11 and re-transmitted to the legitimate host in frame 12.

The usual target will be the subnet's default gateway (the router that accesses other networks). If the ARP poisoning attack is successful, all traffic destined for remote networks will be received by the attacker, implementing an on-path attack.

Domain Name System (DNS) Attacks

The Domain Name System (DNS) resolves requests for named hosts and services to IP addresses. Name resolution is a critical addressing method on the internet and private networks. There are many potential attacks against DNS. On the public internet, attacks might use typosquatting techniques to cause victims to confuse malicious sites with legitimate ones. DNS can be exploited in a DRDoS attack. Threat actors can also directly target public DNS services as a means of performing DoS against a website or cloud resource. Finally, a threat actor might be able to hijack a public DNS server and insert spoofed

records, directing victims to rogue websites.

On a private network, a DNS attack is likely to mean some sort of DNS poisoning. DNS poisoning compromises the process by which clients query name servers to locate the IP address for a domain name. There are several ways that a DNS poisoning attack can be perpetrated.

DNS-Based On-Path Attacks

If the threat actor has access to the same local network as the victim, the attacker can use ARP poisoning to respond to DNS queries from the victim with spoofed replies. This might be combined with a denial-of-service attack on the victim's legitimate DNS server. A rogue DHCP could be used to configure clients with the address of a DNS resolver controlled by the threat actor.

DNS Client Cache Poisoning

Before DNS was developed in the 1980s, name resolution took place using a text file named HOSTS. Each name:IP address mapping was recorded in this file, and systems administrators had to download the latest copy and install it on each internet client or server manually. Even though most name resolution now functions through DNS, the HOSTS file is still present, and most operating systems check the file before using DNS. Its contents are loaded into a cache of known name:IP mappings and the client only contacts a DNS server if the name is not cached. Therefore, if an attacker is able to place a false name:IP address mapping in the HOSTS file and effectively poison the DNS cache, they will be able to redirect traffic. The HOSTS file requires administrator access to modify. In UNIX and Linux systems, it is stored as /etc/hosts, while in Windows, it is placed in %SystemRoot%\System32\Drivers\etc\hosts. The presence of suspect entries in the HOSTS file indicates that the machine has been compromised.

DNS Server Cache Poisoning

DNS server cache poisoning aims to corrupt the records held by the DNS server itself. This can be accomplished by performing DoS against the server that holds the authorized records for the domain and then spoofing replies to requests from other name servers. Another attack involves getting the victim name server to respond to a recursive query from the attacking host. A recursive query compels the DNS server to query the authoritative server for the answer on behalf of the client. The attacker's DNS, masquerading as the authoritative name server, responds with the answer to the query but also includes a lot of false domain:IP mappings for other domains that the victim DNS accepts as genuine. The nslookup or dig tool can be used to query the name records and cached records held by a server to discover whether any false records have been inserted.

DNS Attack Indicators

A DNS server may log an event each time it handles a request to convert between a domain name and an IP address. DNS event logs can hold a variety of information that may supply useful security intelligence and attack indicators, such as:

- The types of queries a host has made to DNS.
- Hosts that are in communication with suspicious IP address ranges or domains.
- Statistical anomalies, such as spikes or consistently large numbers of DNS lookup failures, may point to computers infected with malware, misconfigured, or running obsolete or faulty applications.

DNS is also a popular choice for implementing command & control (C&C) of remote access Trojans. It can be used as a means of covertly exfiltrating data from a private network.

6.5.3 Analyzing ARP Poisoning (Demo Video)

Transcript:

On-path attacks are very popular because they're so productive. When a hacker puts themselves between a victim and the machine they're connecting to, they can capture traffic, delete traffic, inject their own traffic, spoof DNS, and capture passwords. Imagine the damage that can happen if someone captures important documents and leaks their contents or steals the information. In this demo, we're going to illustrate how to capture a password with a on-path attack so that, as a security analyst, you can take countermeasures to protect your organization.

So, let's get started. First, I'm going to scan my subnet to see what hosts are connect and identify my router. Inside this terminal on my Kali Linux machine, I can type in "nmap 10.10.10.0/24' and press Enter. This is going to scan my 10.10.10.0 subnet. It finished pretty quickly.
Now let's scroll up and see what we found. Up there, I found my router, router.CorpNet.xyz with an IP address of 10.10.10.1. We need to remember this information for later.
Here, I have another device on the CorpNet domain with an IP of 10.10.10.169.
I'll scroll down a bit. This device, here, is the one I'm most interested in, PC-12 with an IP ending in 195. This is the machine that's going to be my victim machine and the one that I'm going to try to capture passwords from. Let' go ahead and close this terminal.
Now I want to open Ettercap. Ettercap comes in both a GUI and a command line interface. It's pre-installed as part of Kali Linux, so we don't have to install or download anything. To open Ettercap, I go up here, to Applications, go down to Sniffing & Spoofing, and then, over to the right, find Ettercap. The Ettercap GUI will appear. And from there, I'll come up to the Sniff tab and select Unified Sniffing. That will launch a dialog box, where I'm prompted to select my Network interface. I only have two options. Depending on your system, you may have multiple interfaces such as a WLAN interface or multiple LAN interfaces. I'm going to leave it with eth0 and click OK.
Now let's scan our subnet for hosts. As you'll recall from our NMAP scan, I had three hosts. I'll come up to the Hosts tab, come down to Scan for hosts, and select it. Ettercap starts scanning the subnet for hosts. Down here, you can see that it scanned 255 hosts and found three.
Now I want to list these hosts, so I'll go up to the Host tab. This time, we'll come down to the Hosts list and select it. Here, we can see our three hosts. Remember that this 10.10.10.1 was our router, and our victim machine is the one that ended with 195.
Let's go back up here, to the router, and select it. Down here, I want to select Add to Target 1. Now I'll come back up and select the host ending in 195. And then we'll go back down. This time, we'll select Add to Target 2.
So, what am I doing here? This is where I'm going to insert myself between the router and the victim. To do that, I'm going to go up to the Mitm tab, come down a bit, and select ARP poisoning. In this new dialog box, I want to check the box that says "Sniff remote connections" and click OK.
All right. Now, I'll come up to Start and select Start sniffing. I should have stopped sniffing earlier, but it'll still work fine. Now Ettercap is sniffing our traffic between our victim and the router. We just need to wait for our victim to try to log into a website. Let's jump over to our Windows machine, the one with the IP that ended with 195.
All right. I'm over on the Windows machine now, and I've brought up this website. It's a Moodle site, a popular learning management platform used by many educational organizations. I've hidden the web address, but it's an HTTP site, not HTTPS. Just to confirm that we're on the correct machine, I'll quickly open a command prompt and do an ipconfig to verify the IP address. It's the correct machine. It looks good. I'll close this window.
Let's go ahead and log in with some credentials. I'll just type in the word 'teacher', and I'll use 'letmein' for the password. I'll click Log in. It tells me that I have an invalid login, and it wants me to try logging in again. That's totally okay for what we're demonstrating. Let's jump back over to our Kali Linux machine and see what's happening with Ettercap.
Okay. I'm back on Kali, and here's Ettercap. Down here, at the bottom, you can see that when I attempted to log in, Ettercap captured the username "teacher" and the password "letmein". Of course, this isn't the correct username and password, but Ettercap would capture whatever was entered, and we assume that the user would eventually enter their correct credentials.
The important thing to remember here is that this method does not work with HTTPS. But most users have a habit of using the same credentials for multiple sites. So, once you get the credentials to one site, you'll likely have credentials for multiple sites.
That's it for this demo. In this demo, we used Ettercap to create a on-path attack with ARP poisoning to capture login credentials through a website.

6.5.4 Poison ARP and Analyze with Wireshark (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You believe a hacker has penetrated your network and is using ARP poisoning to infiltrate it.

In this lab, your task is to discover whether ARP poisoning is taking place as follows:

- Use Wireshark to capture packets on the *enp2s0* interface for five seconds.
- Analyze the Wireshark packets to determine whether ARP poisoning is taking place.
- Use the **192.168.0.2** IP address to help make your determination.
- Answer the questions.

Explanation

Complete this lab as follows:

1. Use Wireshark to capture packets on *enp2s0* .
    a. From the Favorites bar, select **Wireshark** .
    b. Maximize the window for better viewing.
    c. Under Capture, select **enp2s0** .
    d. From the menu bar, select the ***blue fin*** to begin a Wireshark capture.
    e. After capturing packets for five seconds, select the ***red box*** to stop the Wireshark capture.
2. Filter for only ARP packets.
    a. In the *Apply a display filter* field, type **arp** and press **Enter** to only show ARP packets.
    b. In the Info column, look for the lines containing the **192.168.0.2** IP address.
3. Answer the questions.
    a. From the top right, select **Answer Questions** .
    b. Answer Questions 1 and 2.
    c. Select **Score Lab** .

6.5.5 Analyzing DNS Poisoning (Demo Video)

Transcript:

DNS Spoofing, sometimes referred to as DNS cache poisoning, is an attack where a rogue device directs DNS requests on a network. This basically means that an attacker manipulates the DNS to redirect DNS traffic. This type of attack is very difficult to detect, which makes it quite dangerous. In this demo, we're going to redirect DNS to our Kali Linux machine, spoof the DNS, and point to a web page I've set up.

For this demo, we're going to use Kali Linux with Ettercap. Ettercap has a GUI interface. That's what we're going to use. But we need to do a few other things before we get started.
I'm going to open up a terminal here, on Kali Linux. We need to check the etter.config file. The easiest way to find this file is to just type 'locate etter.conf' and press Enter. The path I'm interested is this one, right up here. I'm going to use the program Leafpad to open the file. We'll type in 'leafpad /etc/ettercap/etter.conf' and press Enter. And Leafpad opens my file.
Now we need to edit these two lines of code, uid and gid. I need to change both of these values to zero.
After that, we need to scroll down a bit, to where it says Linux. Right under there, we want to find the iptables. These two lines are commented out with the pound sign (#). Just delete the two pound signs from these lines of code. There's one here and one here. So, now that we've done that, we'll save our file and close this window.
Now that we're done with that, we need to edit our etter.dns file. I'm back in the terminal. Once again, to find this file, I can just type in 'locate etter.dns'. And, just like before, I'll use Leafpad to open it by typing 'leafpad etter.dns'. The file opens up.
I need to add a line to this host file. I'll scroll down until I see these few lines with Microsoft in them. Right above those lines, I want to add this line, here: '* A 10.10.10.197'. By doing this, we're telling all of the DNS traffic on the network to go to this IP address, which is actually my Kali Linux system's IP address. The Microsoft lines are not commented out by

default, but I went ahead and put pound signs (#) in front of them. We don't want to redirect their DNS. Before I exit, I want to go up and save the file.
All right. We have some more editing to do here. Now I need to edit my index.html file. Since we're spoofing the DNS and redirecting traffic to our Kali system, I need to give the users something to see, so I'll edit the index page.
To get to that file, I'll come up to Places and come down here, to Computer. From there, I need to find the folder called var and open it. Within var, I'll go over to the www folder, then to the html folder. I've already made some changes in here. One of the things I did was copy my index.html file and rename it. Let's open this old file and see what the contents look like.
Okay, our file is open. And, like I said, by default, this is our index.html file on this Linux system. I didn't want to modify the original in case I need to restore it later, so I made this copy. Now let's close this one and open up the modified copy that I created. Once again, I'll use Leafpad to do that. Here, you can see I made a very simple landing page for our victims. It says, "You have been denied access to this site. Nice try!" I've already saved it, so I'll go ahead and close this file.
We're done editing files. Now we can go to Applications > Sniffing & Spoofing > Ettercap. Ettercap opens for us. And the first thing we need to do here is come up to the Sniff tab and select Unified Sniffing. A dialog box opens up. We want to select our network adapter. In our case, that's eth0. Click OK. We actually want to come back up to Start and select Stop Sniffing because we're not quite ready yet.
We're going to go over to the Hosts tab and select the Scan for Hosts selection. Ettercap will do its thing and start scanning. You can see, down here, that we found three hosts.
Now let's go up to Hosts again. This time, select the Hosts list. And down here, you can see the three hosts. Previously, I ran NMAP and figured out that this top host is my router, and this bottom host is my victim machine, the one I want to attack so that anytime this machine goes to a website, the DNS is going to point them to my Kali Linux machine's IP address. So, while my victim IP is selected, I'll come down here and click Add to Target 1.
Now we'll go to Plugins > Manage the plugins and select the dns_spoof option. Now, if we look down here, in the lower pane, you can see it says, "Activating dns_spoof plugin..."
All right, we're almost done. I need to come up to the MITM tab and select ARP poisoning. When the dialog opens up, we need to check the box for Sniff remote connections and then click OK.
We need to start the Apache web service on the Kali machine to actually make it into a web server. I do that by opening up a terminal and typing in 'service apache2 start'. Nothing really exciting happens, but this command starts the Apache service for us.
That wraps up our steps on this Kali system. Now we need to see what happens when our victim tries to surf the web. We'll do that from a Windows machine.
All right. I'm on a Windows machine now. I'll go to a command prompt to check our IP address. I'll type in 'ipconfig', and we can confirm that the IP ends with 195, so this is the right machine. I'll close this window.
Down here, I'll launch Google Chrome, and I'll type in 'www.testout.com'. When the page launches, it takes me to the index.html page on my Kali Linux system and displays the message we typed. That confirms that our DNS has been spoofed and redirected.
That's it for this demo. In this demo, we configured our Kali Linux machine to redirect DNS requests and send the traffic to our modified index.html page. We discussed how we had to edit the files to make this happen and how to configure Ettercap to perform the DNS spoof.

6.5.6 Poison DNS (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You want to spoof the DNS to redirect traffic as part of a man-in-the-middle attack.

In this lab, your task is to:

- (Optional) From the **Exec** computer, access **rmksupplies.com** and verify that site can be accessed.
- From the Linux **Support** computer, use Ettercap to begin sniffing and scanning for hosts.
- Configure the **Exec** computer (192.168.0.30) as the target 1 machine.
- Initiate DNS spoofing.
- From the **Exec** computer, access **rmksupplies.com** and verify that it has been redirected to a different site.

Explanation

Complete this lab as follows:

1. From the Support computer, use Ettercap to begin sniffing and scanning for hosts.
    a. From the Favorites bar, select **Ettercap** .
    b. Select **Sniff** > **Unified sniffing** .
    c. From the Network Interface drop-down menu, select **enp2s0** .
    d. Select **OK** .
    e. Select **Hosts** > **Scan for hosts** .
2. Configure the Exec computer (192.168.0.30) as the target 1 machine.
    a. Select **Hosts** > **Host list** .
    b. Under IP Address, select **192.168.0.30** .
    c. Select **Add to Target 1** to assign it as the target.
3. Initiate DNS spoofing.
    a. Select **Plugins** > **Manage the plugins** .
    b. Select the **Plugins** tab.
    c. Double-click **dns_spoof** to activate it.
    d. Select **Mitm** > **ARP poisoning** .
    e. Select **Sniff remote connections** and then select **OK** .
4. From the Exec computer, access **rmksupplies.com** .
    a. From the top left, select **Floor 1 Overview** .
    b. Under Executive Office, select **Exec** .
    c. From the taskbar, select **Google Chrome** .
    d. In the URL field, type **rmksupplies.com** and then press **Enter** .
    Notice that the page was redirected to RUS Office Supplies despite the web address staying the same.

6.5.7 Analyzing a SYN Flood (Demo Video)

Transcript:

We've already talked about the three-way handshake, but let's quickly review it again. A three-way handshake starts when a client sends a connection request to a server, saying, "Let's connect." This is called a SYN message. The server acknowledges this request by sending a SYN-ACK message back to the client, saying, "Okay, I'll talk to you." Then the client sends back an ACK message, saying, "Great, we're talking." This is a three-way handshake.

When we're doing a SYN Flood, we're doing only part of this process. The attacker, or the client, doesn't respond back to the final ACK message. The server sits and waits with a half-open connection that's never completed, denying others communication with the server in the meantime. Another method is to spoof the IP address so the server responds to the spoofed IP address, never getting any response.
I'm going to demonstrate how all this works.
First, we need to choose a victim. I've already identified the IP address of my victim on my Kali Linux system, and I've run Zenmap against it. I scanned for all open ports from 0 to 65535. Here, you can see that we have all these ports. We're going to attack port 135. I'll minimize Zenmap for now.
To get an idea of what happens during this attack, let's open up the victim machine. It's a Windows 10 machine, and I've installed Wireshark on it. I'm going to go ahead and open Wireshark by clicking on the shark fin. We're doing this to see what happens when this machine is attacked.
I'm also going to come open Task Manager. We need the Performance tab. You can see my CPU is relaxed. We had a spike here, probably when we opened Wireshark. Down here, we really have no Ethernet activity. I actually need to go back over to Wireshark and click the shark fin to start sniffing traffic. I'll just leave for now and go back to Kali Linux.
I'm in Kali Linux, and we're going to use Metasploit for this attack. First, I'll open a terminal here. After it opens, I want to start the database. The Kali Linux Network Services Policy states that no network services, including database services, will run on boot by default, so we need to get the database going in order to get Metasploit up and running with database

support. From the terminal, we'll type 'service postgresql start' and press Enter. This starts the database that we'll need. A lot of times, I'll forget to do this, and you'll get an error letting you know it can't connect to the database.
All right. Now, let's start Metasploit. Type in 'msfconsole' and press Enter. This sometimes takes a few minutes to get started, so be patient. I'll adjust the terminal window and clear the screen.
Now, with Metasploit up and running, I want to search for SYN flood. I'll type in 'search synflood' and press Enter. And I get a matching module.
This is easy because there's only one module listed, so that's the one we're going to use. I'll use that module by copying this part, here. We'll type in 'use', paste in auxiliary/dos/tcp/synflood, and press Enter.
Now you can see that it loaded that module here, in red. We want to see the options for using this, so let's type 'show options'.
Okay. There are a few things we need to do here. We need to configure the remote host, shown here. I also want to configure the remote port. If you remember from our Zenmap scan, that's going to be port 135. I'm also going to spoof my IP address, down here. Okay. Let's configure these.
I'll start with the remote target. I'll type in 'set RHOST 10.10.10.195' and press Enter. Let's set the remote port, 'set RPORT 135', and press Enter.
Next, let's set the spoofed IP address by typing 'set SHOST 10.10.10.50'. I just randomly picked that 50, but I do know that this IP isn't used on my network.
Now, let's look at our options again by typing 'show options'. We have everything in there, and it's just the way we planned, so it looks like we're ready to attack. I'll type 'clear' to clean up the screen.
To run the exploit and start our SYN Flood, we just type in 'exploit' and press Enter. It says, "Hey, I'm SYN flooding 10.10.10.195 on port 135." Well, to be sure, let's jump back over to our Windows 10, the victim, and see what's going on.
I'm on the Windows 10 machine. First, let's look at Wireshark. I have a ton of traffic going from 10.50 to 10.195. Over here, you can see that SYN requests are happening. I'm going to stop sniffing here. Now let's highlight one of these. If I look in the lower pane, I can see one thing, the source's MAC address. It ends in 1c. That can help me determine the attacker's real IP address, since the MAC isn't spoofed. Now let's look at Task Manger.
As you recall, we looked at the CPU and Ethernet usage here. You can see that my CPU is a little higher than what it was previously; but it's not as high as I would have guessed. Also, if you look at the Ethernet, you can see that it's really busy. Remember, that's an attack from just one machine. Imagine if it were coming from 10 machines or hundreds of machines. It wouldn't take much to overwhelm this system.
So, how do you stop SYN floods from happening? There are several methods. Admins used to block the port of the incoming SYN requests, but that didn't work very well. Now, admins can use SYN cookies, RST cookies, filtering, and Microblocks to mitigate SYN flood attacks, as well as other methods.
That's it for this demo. In this demo, we started by reviewing the three-way TCP handshake. We discussed how a SYN flood begins with an incomplete handshake, leaving the server sitting and waiting. We performed a SYN flood using Metasploit. We viewed Wireshark and our performance on the victim machine to analyze the attack's effect. We ended by talking about a few of the countermeasures that help prevent SYN flooding.

6.5.8 Analyze a SYN Flood Attack (Simulation)

Scenario

You are the CorpNet IT administrator. Your support team says that CorpNet's customers are unable to browse to the public-facing web server. You suspect it might be under a denial-of-service attack, possibly a TCP-SYN flood attack. Your www_stage computer is on the same network segment as your web server, so you should use this computer to investigate the problem.

In this lab, your task is to:

- Capture packets from the network segment on **www_stage** using Wireshark.
  - Use the **enp2s0** interface.
- Analyze the attack using the following filters:
  - **tcp.flags.syn==1 and tcp.flags.ack==1**
  - **tcp.flags.syn==1 and tcp.flags.ack==0**
- Answer the question.

Explanation

Complete this lab as follows:

1. Using Wireshark, only capture packets containing both the SYN flag and ACK flags.
    a. From the Favorites bar, select **Wireshark** .
    b. Under Capture, select **enp2s0** .
    c. From the menu, select the ***blue fin*** to begin the capture.
    d. In the *Apply a display filter* field, type **tcp.flags.syn==1 and tcp.flags.ack==1** and press **Enter** to filter Wireshark to display only those packets with the SYN flag and ACK flag. You may have to wait up to a minute before any SYN-ACK packets are captured and displayed.
    e. Select the **red square** to stop the capture.
2. Change the filter to only display packets with the SYN flag.
    a. In the *Apply a display filter* field, change the tcp.flags.ack ending from the number **1** to the number **0** and press **Enter** .
    Notice that there is a flood of SYN packets being sent to 128.28.1.1 (www.corpnet.local) that are not being acknowledged.
    b. From the top right, select **Answer Questions** .
    c. Answer the question.
    d. Select **Score Lab** .

6.5.9 Examining DNS Attacks (Demo Video)

Transcript:

In this demo, we're going to look at a few ways an attacker may use the DNS system to get information about the computers in an organization, and also to compromise a machine and get access to information.

Let's first look at doing reconnaissance using DNS servers. One way this can be done is with the nslookup command.
You see that I have a PowerShell command prompt window open here, and if you just type nslookup, that will put you into this utility. nslookup is used to troubleshoot DNS server settings.
You can see here that the DNS server that I'll be using is this one. It's the default one that I have configured for my machine here. When a computer wants to resolve a name, the first thing it does is look into a special file called a host file.
We'll talk more about that later, but essentially the host file can contain mappings between domain names and specific IP addresses. If it's not in the host file, then an operating system checks its DNS cache.
If the mapping is in there, then it uses the IP address that it finds. If not, it goes out and does the name resolution with the DNS server that's configured.
In this case, nslookup simply allows us to do those queries by passing the cache that might be on a local operating system.
So, if we want to, here, we can simply type in the name that we want to resolve like google.com, and it will go out and resolve that. My DNS server here goes out and says, I don't know what IP address that is associated with.
And then it issues the query to the resolver that it's pointing to. And you see here the IP address of www.google.com. Of course, we can try other ones like testout.com and others in you'll see that the resolution occurs appropriately.
For this demo, I've set up a domain that we can use to see how a few vulnerabilities work.
One vulnerability that exists with DNS servers is the ability to find out all the domain information associated with a particular name. One way we can do that is using the list command.
To do that, we type ls; we specify the domain that we want using the -d flag; and then we can say ajptech.net. This will tell us all of the computers that are associated with the ajptech domain.
If you do that, you see that we have a full list of all the machines.
Now, this is just a test environment. I've just listed a few servers, etc.
You can imagine in an organization that has hundreds of computers, this provides very nice information to know which IP addresses we should spend time attacking, based on these records.

This is referred to as a zone transfer, and we need to make sure that our DNS servers in our organization don't allow zone transfers to just anyone.
They're useful because you don't want just a single server, and so you want to share those records with other secondary zones, but you need to restrict who is allowed to do those zone transfers.
Let's take a look at that, inside of Windows Server 2019. You see here I have launched the DNS manager on the server. If you haven't installed that role before, it's super simple to do, and you do so before configuring these different zones.
Let's go ahead and look at some of the properties that you might want to set.
The one that's of most interest to us right now is this Zone Transfers tab. This basically specifies who is allowed to gather all those records at once, and you can see here that we have 'To any server'.
We could come in and just say, 'No zone transfers are allowed.' But as I explained a minute ago, that's useful, because if this is the primary zone for that domain name, then you want to have a secondary zone.
You don't want a single point of failure here, and so you want to be able to allow zone transfers.
You can only allow servers in the Name Servers tab, which is nice because usually you set those up specific to a domain, and that would be anything that's listed here, or you can restrict it to a specific IP address.
Let's go ahead and apply those settings. Let's say that we're going to only allow the following servers, and I'll just put in a fake IP address here right now. It won't resolve right now, but that's okay. But it will restrict those zone transfers.
Let's go ahead and switch back over to our PowerShell window, and let's try that same command. You can see now that it says it has refused transferring all that information from the zone--perfect.
Now when people are doing reconnaissance, they won't be able to see everything that we have in that zone.
However, the domain name server continues to work, so I can say server1, and you can see that that name resolution continues to work. So, make sure you secure your zone transfers.
Today, restricting zone transfers is commonplace; in fact, it's the default of many operating systems, and it's the best practice everybody should have. If your DNS server isn't configured that way, you want to make sure that you take care of that. Now let's move on to a couple of other concepts.
There are many different DNS attacks, such as distributor reflection, denial of service attacks, cache poisoning, TCP send floods, and others.
Let me just briefly explain one of these to give you a flavor.
A distributor reflection denial of service attack occurs when several DNS queries are sent to a DNS server that has an open resolver, and that DNS server responds to the spoofed address with a large payload.
The key to the attack is that the attacker only has to send a small packet of information to the DNS server, and then the DNS server sends a large response, but that response doesn't go back to the attacker; it sends it to the target that the attacker listed, using a spoofed source IP address.
In short, the attacker uses the DNS to create a denial of service attack against the third party. In this case, you need to make sure that you don't have any open resolvers.
An open resolver is essentially a DNS server that's willing to resolve names for third parties.
Usually you protect this in a couple of different ways--either restricting the DNS requests that come into your DNS server, that is, you don't allow outside requests; or simply not having any forwarding capabilities associated with your DNS server, and making sure that recursion is disabled on your DNS server.
Let me show you that setting quickly here. If you come to your server, you'll see the properties of the server itself. We are in the zone from before. We're going to go over here to the Advanced tab.
You want to make sure that this option is set, so that you disable the recursion. That, of course, disables the forwarders as indicated there. I'll say OK there.
Then I'll come over to the Forwarders tab here. You can see that I'm set up, but it's not using those anymore. The downside of this is we need to use a different DNS server in order to do our name resolution.
This server now can only do name resolution for the domains that are on it, but that's okay. That's best practice. You definitely don't want to be susceptible to those amplification attacks that I just explained.
Let's come back here to our PowerShell prompt and get out of the nslookup command.
Let me show you a couple of other things. I'll go ahead and clear the screen, as well.
A minute ago, I explained that when name resolution needs to occur on an operating system, it first checks the set host file, and then it goes to cache, sees what's in cache and uses that if there is something that's associated with that name, and then it goes and does the name resolution using the DNS system.
Let's go ahead and take a look at that cache.
But before we do so, we need to make sure that we have something in cache. I'll simply go ahead and try to ping google.com. You can see that I get a response.
Now, to see what's in cache, you simply say ipconfig /displaydns. You can see that I have several different entries in there.

I have some of those entries because I have a browser running on my system, and it's already resolved some names.
Let's go ahead and try to clear that out. To clear out your cache, you type ipconfig /flushdns.
Let's ping google.com, and then let's go ahead and see if we can see just those name resolutions.
You can see here there is only one record now when we display the dns cache.
So, this cache constantly keeps track of what's been resolved previously so it doesn't have to invoke the DNS system for every single request if it already knows that information. It saves time.
Each of these records also has a time to live, as you can see there, and so they're constantly refreshing as the system is being used.
One mistake that some people come across when they're trying to resolve DNS issues is they forget to clear out their cache.
They hit a site that's in cache, they make some changes, and they don't realize that it's actually pulling it out of cache instead of going on. I want to make sure you're aware of how cache works.
Now what if we could poison the cache or insert a record in there that wasn't legitimate, that came from another source?
This is a little bit harder to architect, but if we could specify a different IP address for the name server, the DNS server, we could pull records from it. It would be saved in the cache or locally, and then subsequent resolutions would pull the IP address out of the cache.
Indeed, that's one attack that individuals use when exploiting the DNS system. We're not going to show that to you today, but I want to show you one other concept before we wrap up this demo, and that has to do with the host file that I mentioned a few minutes ago.
This is the host file I mentioned a minute ago, and you can see that it's located in Windows System32\drivers\etc\hosts.
In this file, you can include specific mappings between a domain name and an IP address.
Before there was this DNS system, this file was distributed to all the computers on the internet, and so there's just a huge file of these mappings so that you wouldn't have to remember the specific IP addresses.
Thankfully, the DNS system came along, making our lives a lot easier and being allowed to update, but this is still in place. And as I mentioned, it gets checked before the OS resolver cache, before going out to the DNS system to resolve domain names.
You can see an example of how you can use this. But before I put in an example, let me show you a legitimate site so that you can see how this works. Let's go to cat-bounce.com, just a random site.
You can see that there's something out there. If we want to redirect individuals, instead of them going to this site to a different site, we can go into the host file and add a mapping that we're interested in.
You can see that we put the IP address first, and then the domain name. So, I know that the duckduckgo server is at the following address, and then I can simply put www.cat-bounce.com, and I'll save that.
Recent browsers have incorporated a DNS cache inside of the browser itself.
They do this because it gives them a competitive advantage to go ahead and resolve those domain names ahead of time, the minute that they show up on a page and that you might potentially click on.
Lets open an incognito window to bypass the cache.
If we go to www.cat-bounce.com, you see that it forwards on duckduckgo.com. I hope you can see how dangerous this could be. This is considered host file manipulation.
However, an attacker could simply keep that in place and use. An innocent individual wouldn't have any idea that you weren't on the legitimate site.
Consider banking sites or corporate sites, anywhere where sensitive information might be shared. An attacker can get access to this host file, present what looks to be a legitimate website, and then siphon off data as you use it.
It's also good to know that you can use host file to block malicious information, if you are a systems administrator.
This is a common site. It looks a little dated, but the information is pretty good on it. You can come and download a specific host file to block out a bunch of bad or malicious sites.
The IP address you see is a null IP address. In other words, if we added this to our host file and you went there, it simply wouldn't go anywhere.
One thing to note, you'll need to have administrator rights to adjust the host file.
That's it for this demo. In this demo, we've look at a few ways that an attacker can use the DNS system to get out information about computers in the organization and to compromise a machine and eventually grab sensitive information.
We've looked at the nslookup command. We've looked at zone transfers. We've looked at securing the open resolvers.
And we've looked at the host file and clearing cache.
It's really important that you protect your DNS servers and make sure that you aren't taken advantage of in that way.

6.5.10 Malicious Code (Lesson Video)

Transcript:

Malicious software is perhaps the most dangerous threat to any computing device. It's so prevalent that most people combine the words â€˜malicious software' and just call it malware. It isn't always coded using advanced programming languages. In fact, Python is one of the common languages used to design malicious codes. Sometimes a simple script can exploit built-in OS features such as shells. Even application features like macros can be used to develop and execute malicious codes. In this lesson, we'll look at how malicious code and scripts are designed and executed.

Python is one of the most popular programming languages. It was first released in 1999. It's designed to be easy to learn and read, and it works with most operating systems, including Windows, MacOS, and Linux.
Aside from being easy to read and understand, Python can take advantage of a massive amount of open-source Python packages and repositories. These add-ons make it that much easier to develop a Python script to do whatever you want.
Python is often used in the development of Remote Access Trojans, or RATs. This is because Python makes it very easy to implement libraries that can take screenshots, use a webcam, and even make web requests. It's also easy develop malicious code that can run on many different systems and devices, including Android devices.
One of the main drawbacks to using Python for malware is that its file sizes are larger than other common languages. Python must also be installed on the system for a Python script to run. This is fine for MacOS and Linux. Windows doesn't come with it installed, but it's pretty easy to convert Python scripts to Windows-compatible executables.
In addition to Python, attackers often use shells. A shell is a program that provides an interface that gives users access to operating system functions and services. We generally associate shells with command line interfaces, but they can have graphical interfaces too. These programs can provide access to core operating system functions, which makes them exploitation targets.
You can type commands directly into the shell window or run pre-written scripts. A script is simply a plain text file with the commands you want to run typed in, just like you would type them into the shell window. When the script runs, the commands are executed.
Now let's talk about PowerShell. PowerShell is a management framework that Microsoft developed to replace the Command Prompt and give users more power and control over the Windows system. PowerShell is built on the .NET frameworks and can run on multiple operating systems, including MacOS and Linux.
PowerShell uses cmdlets to carry out commands. Cmdlets are tiny scripts that are used to perform certain functions. Some cmdlets replace older commands and provide more advanced functions. Users can combine these cmdlets to develop scripts that automate tasks and configure just about anything in Windows.
Malicious PowerShell scripts pose a major security threat. For example, let's talk about fileless malware. PowerShell scripts can run in the memory of the system--they don't need an actual file executable to run. An attacker can take advantage of this function by sending a phishing email. When the victim clicks the link, a PowerShell script is loaded in the background that can download and run other malware programs. Because these attacks are fileless, anti-virus programs generally can't detect them.
Bash is a command shell and scripting language used in most Linux distros and MacOS versions prior to Catalina. Bash was released in 1989, and it's still heavily used. When you use a command in Linux, Bash works in the background to execute that command using environment variables.
Since many web servers run on Linux's Apache platform, Bash malware can be designed to attack these systems. A well-known malware called Shellshock uses Bash commands to exploit a flaw within the Bash shell itself, allowing an attacker to inject malicious commands.
Malware that exploits these shell programs essentially use the operating system against itself. This makes these scripts difficult to detect and prevent. Keeping anti-malware up to date and training users not to click unknown links or run unknown programs is vital.
All right. Next, we're going to talk about macros, which are similar to scripts. Macros are little bits of code that can be used to run a series of steps or functions in an application, making common tasks more efficient. Many different programs can make use of macros, but the most common use of macros is in the Microsoft Office programs.
Microsoft Office programs can use the Visual Basic for Applications, or VBA, programming language to create and run macros. VBA can be a powerful programming language and can even be used to run commands in Windows using the shell function or delete files from your hard drive using the Kill function.
An example of a devastating macro malware is the macro virus Melissa. In 1999, Melissa was distributed as a Word document, and when it was opened, a macro ran that pulled the first 50 users from the user's Outlook address book and mailed a copy of the infected Word document. Melissa spread extremely quickly and shut down more than 300 corporations and government agencies temporarily as their networks were overloaded. The total estimated damage of Melissa was around $80 million dollars.

In newer version of Microsoft Office, macros are disabled by default, and a user must specifically allow them to run. This helps to keep our systems safe from malicious code. But if someone uses macros often and has them enabled, their system is extremely vulnerable.
That'll wrap up this lesson. We've covered some methods for generating and distributing malicious code. First, we looked at the popular scripting language Python. Python is used to create all sorts of great programs and scripts, but because of its ease of use, many hackers use it to develop malicious code.
Then we covered command shells like PowerShell and Bash. These shells generally allow access to all sorts of critical operating system functions, meaning malicious code can cause them serious damage.
Finally, we looked at how macros can be exploited to run malicious code on a computer. Microsoft Office programs use the VBA programming language to develop macros to automate tasks, but it can be exploited to distribute malicious code.

6.5.11 Malicious Code Facts

Malicious software (malware) is perhaps the most dangerous threat to any computing device. Malware can be created using a variety of programming languages and methods.

This lesson covers the following topics:

- Malicious code indicators
- Python
- Command shells
- Macros

Malicious Code Indicators

Many network attacks are launched by compromised hosts running various types of malicious code. Indicators of malicious code execution are either caught by endpoint protection software or discovered after the fact in logs of how the malware interacted with the network, file system, and registry. To understand how and where these indicators are generated, it is helpful to consider the main types of malicious activity:

- Shellcode — this is a minimal program designed to exploit a vulnerability in the OS or a legitimate app to gain privileges or to drop a backdoor on the host if run as a Trojan. Having gained a foothold, this type of attack will be followed by some type of network connection to download additional tools.
- Credential dumping — the malware might try to access the credentials file (SAM on a local Windows workstation) or sniff credentials held in memory by the lsass.exe system process. Additionally, a DCSync attack attempts to trick a domain controller into replicating its user list along with its credentials with a rogue host.
- Pivoting/lateral movement/insider attack — the general procedure is to use the foothold to execute a process remotely, using a tool such as PsExec or PowerShell.The attacker might be seeking data assets or may try to widen access by changing the system security configuration, such as opening a firewall port or creating an account. If the attacker has compromised an account, these commands can blend in with ordinary network operations, though they could be anomalous behavior for that account.
- Persistence — this is a mechanism that allows the threat actor's backdoor to restart if the host reboots or the user logs off. Typical methods are to use AutoRun keys in the registry, add a scheduled task, or use Windows Management Instrumentation (WMI) event subscriptions.

Python

Python has become one of the most popular programming languages. First released in 1991, Python is designed to be easy to learn and read. It can be used on most operating systems, including Windows, MacOS, and Linux. Python can also take advantage of open-source Python packages and repositories.

Many Remote Access Trojans (RATs) are designed using Python. Python makes it easy to implement libraries that allow the RAT to perform functions such as:

- Taking screenshots.
- Enabling the webcam and viewing it remotely.
- Making web requests.
- Making phone calls.

Python also makes it very simple to develop malicious code that can be run on many different systems and devices, including Android devices.

One of the main drawbacks to using Python for malware is the file size. Python files are larger than other common languages. Also, Python must be installed on a system for a Python script to run. This works for MacOS and Linux, but Windows does not come with it installed. Python scripts can be converted to Windows-compatible executables fairly easily.

Command Shells

A shell provides an interface for users to access operating system functions and services. Shells are generally associated with command line interfaces, but they can have graphical interfaces also. Because these programs provide access to core operating system functions, they are extremely dangerous when exploited.

Commands can be typed directly into the shell program or can be run from a script. A script is a plain-text document that has the commands typed out just like they would be in the shell. When the script is run, the commands are executed.

Two of the more heavily used shells are PowerShell and Bash. The following table describes these two shell programs:

Shell Program	Description
PowerShell	PowerShell is a management framework that Microsoft developed to replace Command Prompt and give users more power and control over the Windows system. PowerShell is built on the .NET framework and can now be run on multiple operating systems, including MacOS and Linux. PowerShell uses cmdlets to execute commands. Cmdlets are tiny scripts that perform certain functions. Some cmdlets replace older commands and provide more advanced functions. Users can combine these cmdlets to develop scripts to automate tasks and configure just about anything in Windows. Malicious PowerShell scripts pose a major security threat. These scripts can run in the memory of the system, which means they do not need an executable to run. • An attacker can take advantage by running malicious PowerShell scripts in the background. • This type of malware is known as *fileless malware.* Fileless malware is especially dangerous because many anti-virus programs are unable to detect it.
Bash	Bash is a command shell and scripting language used in most Linux distributions and MacOS versions prior to Catalina.

	Bash was released in 1989 and is still heavily used. When a command is executed in Linux, Bash works in the background to execute the command using environment variables. Since many web servers run on Linux's Apache platform, malware can be designed in Bash to attack these systems. A well-known malware called Shellshock uses Bash commands to exploit a flaw within the Bash shell. The flaw allows an attacker to inject malicious commands.

Malware that exploits shell programs uses the operating system against itself, making this type of malware difficult to detect and prevent.

Keeping anti-malware up to date and training users not to click unknown links or run unknown programs is vital to staying safe.

Macros

Macros are similar to scripts in that they are little bits of code used to perform a series of steps or functions. Macros, however, are used inside specific applications. Many different programs can make use of macros, but the most common use of macros is in the Microsoft Office programs.

Microsoft Office programs use the Visual Basic for Applications (VBA) programming language to create and run macros. If the Office program is not configured properly, malicious VBA code can be used to open a shell on the Windows operating system. The shell can be used to perform malicious attacks.

In newer versions of Microsoft Office, macros are disabled by default, and a user must specifically allow them to run.

6.5.12 Practice Questions (Section Quiz)

q_analyz_netattacks_amplified_secp8

A threat actor has initiated an attack against a company's network. The threat actor is exploiting a weakness in a specific application protocol, manipulating the request in such a way that the target is forced to respond with a large amount of data, thereby consuming a significant amount of the target's bandwidth.

Which type of attack does this scenario represent?

Answers:

- Distributed denial-of-service (DDoS) attack
- Reflected attack
- ***Amplified attack**
- Direct attack

Explanation:

Amplified attack is the correct answer. An amplification attack targets weaknesses in specific application protocols to make the attack more effective at consuming target bandwidth. The threat actor manipulates the request in such a way that the target is forced to respond with a large amount of data, which is exactly what is described in the scenario.

Distributed denial-of-service (DDoS) attack is incorrect because while a DDoS attack involves multiple hosts launching an attack simultaneously, it does not specifically involve the exploitation of a weakness in a specific application protocol to force the target to respond with a large amount of data.

Reflected attack is incorrect because a reflected attack involves the threat actor spoofing the victim's IP address and using third-party servers to direct responses to the victim host. The scenario does not mention the use of third-party servers or IP address spoofing.

Direct attack is incorrect because a direct attack involves the threat actor directly attacking the victim's network without the use of third-party servers or specific application protocols. The scenario describes the exploitation of a specific application protocol, which is characteristic of an amplified attack.

q_analyz_netattacks_arp_pois_01_secp8

Which of the following is the term used to describe what happens when an attacker sends falsified messages to link their MAC address with the IP address of a legitimate computer or server on a network?

Answers:

- Port mirroring
- MAC flooding
- ***ARP poisoning**
- MAC spoofing

Explanation:

Address Resolution Protocol (ARP) poisoning is when an attacker sends fake ARP messages to link their MAC address with the IP address of a legitimate computer or server on the network. Once their MAC address is linked to an authentic IP address, the attacker can receive any messages directed to the legitimate address. As a result, the attacker can intercept, modify, or block communications to the legitimate MAC address.

Port mirroring creates a duplicate of all network traffic on a port and sends it to another device.

MAC flooding is when an attacker intentionally floods a content-addressable memory table with Ethernet frames, each originating from different MAC addresses. Once the table starts to overflow, the switch responds by broadcasting all incoming data to all ports, basically turning itself into a hub instead of a switch.

MAC spoofing is done to enable the bypass of access control lists on servers or routers by either hiding a computer on a network or by allowing it to impersonate another network device.

q_analyz_netattacks_arp_pois_02_secp8

Which of the following attacks tries to associate an incorrect MAC address with a known IP address?

Answers:

- ***ARP poisoning**
- Hijacking
- Null session
- MAC flooding

Explanation:

ARP spoofing/poisoning associates the attacker's MAC address with the IP address of a victim's device. When computers send an ARP request to get the MAC address of a known IP address, the attacker's system responds with its MAC address.

MAC flooding overloads the switch's MAC forwarding table to make the switch function like a hub. The attacker floods the switch with packets, each containing different source MAC addresses. The flood of packets fills up the forwarding table and consumes so much of the memory in the switch that it causes the switch to enter a state called Fail Open mode, in which all incoming packets are broadcast out of all ports (as with a hub), instead of just to the correct ports.

A null session is the ability to log on using a blank username and password. With hijacking, an attacker steals an open session, inserting themself into the session in place of the original client.

q_analyz_netattacks_client_cache_poisoning_secp8

A threat actor has successfully manipulated a client's DNS cache, causing the client to resolve domain names to incorrect IP addresses controlled by the threat actor. This allows the threat actor to redirect the client's network traffic to malicious websites.

Which type of attack does this scenario represent?

Answers:

- Man-in-the-middle attack
- On-path attack
- ***Client cache poisoning**
- Distributed denial-of-service (DDoS)

Explanation:

Client cache poisoning is the correct answer. Client cache poisoning involves a threat actor manipulating a client's DNS cache, causing the client to resolve domain names to incorrect IP addresses. This is exactly what is described in the scenario.

While a man-in-the-middle attack does involve intercepting and possibly altering communication between two parties, it typically involves the threat actor actively impersonating one or both parties to deceive the other. The scenario does not mention any impersonation, only manipulation of the client's DNS cache.

On-path attack is incorrect because it involves a threat actor positioning themselves in the communication path between two parties and intercepting and possibly altering the communication without their knowledge. The scenario does not describe any interception of communication, only manipulation of the client's DNS cache.

Distributed denial-of-service (DDoS) attack is incorrect because a DDoS attack involves overwhelming a network or service with traffic from multiple sources to cause a denial-of-service. The scenario does not describe any overwhelming of network resources or denial of service.

q_analyz_netattacks_ddos_secp8

A major online retail company has recently been experiencing intermittent downtime of its website. Network analysts observe a massive influx of traffic from multiple sources to the server. However, the traffic seems redirected from other systems.

What type of attack is the company likely experiencing?

Answers:

- ***Distributed denial-of-service (DDoS)**

- Collision
- Injection
- Buffer overflow

Explanation:

In this scenario, distributed denial-of-service (DDoS) attacks align with the influx of traffic from multiple sources, causing downtime.

Collision is mainly used in the context of cryptographic hash functions and does not relate to the influx of traffic causing downtime.

Injection attacks occur when an attacker sends harmful data to a system, tricking it into executing unintended commands. They typically do not result in a massive influx of traffic.

Buffer overflow can lead to crashes or exploitable conditions, but it is not a direct cause of a flood of internet traffic.

q_analyz_netattacks_dns_pois_01_secp8

Which type of denial-of-service (DoS) attack occurs when a name server receives malicious or misleading data that incorrectly maps host names and IP addresses?

Answers:

- ***DNS poisoning**
- ARP poisoning
- Spam
- SYN flood

Explanation:

DNS poisoning occurs when a name server receives malicious or misleading data that incorrectly maps host names and IP addresses. In a DNS poisoning attack:

- Incorrect DNS data is introduced into a primary DNS server.
- The incorrect mapping is made available to client applications through the resolver.
- Traffic is directed to incorrect sites.

ARP poisoning corrupts the ARP cache or sends incorrect ARP data that spoofs MAC addresses, causing devices to send frames to the wrong host or an unreachable host.

Spam sent in such great amounts can consume bandwidth or fill a mailbox, leaving no room for legitimate traffic.

A SYN flood exploits the TCP three-way handshake.

q_analyz_netattacks_dns_pois_02_secp8

While using the internet, you type the URL of one of your favorite sites in the browser. Instead of going to the correct site, the browser displays a completely different website. When you use the IP address of the web server, the correct site is displayed.

Which type of attack has likely occurred?

Answers:

- ***DNS poisoning**
- Hijacking
- Man-in-the-middle
- Spoofing

Explanation:

Because the correct site shows when you use the IP address, you know that the main website is still functional and that the problem is likely caused by an incorrect domain name mapping. DNS poisoning occurs when a name server receives malicious or misleading data that incorrectly maps host names and IP addresses. In a DNS poisoning attack:

- Incorrect DNS data is introduced into the cache of a primary DNS server.
- The incorrect mapping is made available to client applications through the resolver.

Spoofing is used to hide the true source of packets or redirect traffic to another location. Spoofing attacks use modified source and destination addresses in packets and can include site spoofing that tricks users into revealing information.

A man-in-the-middle attack is used to intercept information passing between two communication partners. TCP/IP hijacking is an extension of a man-in-the-middle attack in which the attacker steals an open and active communication session from a legitimate user.

With spoofing, man-in-the-middle, and hijacking, the attack would be successful regardless of whether the DNS name or the IP address was used.

q_analyz_netattacks_dns_pois_03_secp8

An attacker uses an exploit to push a modified hosts file to client systems. This hosts file redirects traffic from legitimate tax preparation sites to malicious sites to gather personal and financial information.

Which kind of exploit has been used in this scenario?

Answers:

- Reconnaissance
- ***DNS poisoning**
- Man-in-the-middle
- Domain name kiting

Explanation:

DNS poisoning (also known as DNS cache poisoning) occurs when a name server receives malicious or misleading data that incorrectly maps host names and IP addresses. In a DNS poisoning attack:

- Incorrect DNS data is introduced into the cache of a primary DNS server.
- The incorrect mapping is made available to client applications.

Reconnaissance is used to gather information for an attack. The goal is to obtain DNS records that identify computer names and IP addresses in a network.

Domain name kiting occurs when spammers exploit domain registration by taking advantage of the five-day grace period for a newly registered domain name to acquire domains and never pay for the registration of domain names. They accomplish this by unregistering a domain name just before the grace period is up and then immediately re-registering the domain name.

Man-in-the-middle attacks are used to intercept information passing between two communication partners.

q_analyz_netattacks_dns_pois_04_secp8

A company CEO is upset after receiving a call from a reporter at a local news station that the company is apparently at a launching point for a massive attack. The reporter provided detailed IP logs, and the network team reviewed them but could not find similar entries.

What could be a possible explanation for the different records?

Answers:

- ***DNS poisoning**
- DNS client cache poisoning
- DNS-based on-path attack
- DNS attack indicators

Explanation:

Domain Name System (DNS) poisoning is an attack in which network traffic appears to be coming from another location to confuse the company looking for an infection that does not exist.

DNS client cache poisoning occurs when an attacker accesses the local client's host file, placing a false name-to-IP address mapping in the file, effectively poisoning the DNS cache to redirect traffic.

Within a DNS-based on-path attack, if the threat actor has access to the same local network, the attacker can use Address Resolution Protocol (ARP) poisoning to respond to DNS queries from the victim with spoofed replies.

DNS attack indicators are bits of information from a DNS server to assist with security intelligence and attack indicators.

q_analyz_netattacks_lifecycle_01_secp8

A threat actor has successfully breached a company's network and has installed malicious code on a compromised host. The threat actor is now operating the compromised host remotely and maintaining access to it over a period of time. The threat actor's activity is disguised as part of the network's regular traffic.

Detection of this type of activity usually depends on identifying anomalous connection endpoints.

Which stage of the cyberattack lifecycle does this scenario represent?

Answers:

- Reconnaissance
- Weaponization, delivery, and breach
- ***Command and Control**
- Data exfiltration

Explanation:

Command and Control is the correct answer. In this stage, the threat actor operates a compromised host remotely and maintains access to it over a period of time. The threat actor has to disguise the incoming command and outgoing beaconing activity as part of the network's regular traffic, which is exactly what is described in the scenario.

Reconnaissance is incorrect because reconnaissance refers to the initial stage of the cyberattack lifecycle where the threat actor uses scanning tools to learn about the network. In the given scenario, the threat actor has already breached the network and is operating a compromised host remotely.

Weaponization, delivery, and breach are incorrect because this stage refers to techniques that allow a threat actor to get access without having to authenticate. This typically involves various types of malicious code being directed at a vulnerable application host or service over the network. In the given scenario, the threat actor has already gained access and is now maintaining it.

Data exfiltration is incorrect because data exfiltration refers to obtaining an information asset and copying it to the attacker's remote machine. In the given scenario, there is no mention of the threat actor obtaining or copying any information.

q_analyz_netattacks_lifecycle_02_secp8.

A threat actor is attempting to learn passwords or cryptographic secrets that will allow them to obtain authenticated access to network systems.

They are using various techniques to gather this information without yet having breached the network.

Which stage of the cyberattack lifecycle does this scenario represent?

Answers:

- Reconnaissance
- ***Credential harvesting**
- Command and Control
- Data exfiltration

Explanation:

Credential harvesting is the correct answer. In this stage, the threat actor attempts to learn passwords or cryptographic secrets that will allow them to obtain authenticated access to network systems. This is exactly what is described in the scenario.

Reconnaissance is incorrect because reconnaissance refers to the initial stage of the cyberattack lifecycle where the threat actor uses scanning tools to learn about the network. While this stage may involve some gathering of information, it does not specifically involve the collection of passwords or cryptographic secrets.

Command and Control is incorrect because in this stage, the threat actor operates a compromised host remotely and maintains access to it over a period of time. The scenario does not describe the threat actor having breached the network or operating any hosts remotely.

Data exfiltration is incorrect because data exfiltration refers to obtaining an information asset and copying it to the attacker's remote machine. In the given scenario, there is no mention of the threat actor obtaining or copying any information.

q_analyz_netattacks_on-path_01_secp8

A network administrator suspects an attacker is intercepting and potentially modifying communications between their organization's server and the client systems. The attacker is not detected by either party during this process.

Which type of attack is the network administrator likely observing in this instance?

Answers:

- ***On-path attack**
- Domain Name System (DNS) attack
- Replay attack
- Distributed denial-of-service (DDoS) attack

Explanation:

An on-path attack involves an attacker intercepting and potentially modifying communications between two parties who believe they're communicating directly with each other.

A DNS attack involves the manipulation of DNS entries to redirect traffic to malicious sites. While it is a severe threat, it does not include the direct interception or modification of ongoing communications.

A replay attack involves the interception and retransmission of a valid data transmission. It does not include the real-time interception and modification of communications between two parties.

A distributed denial-of-service (DDoS) attack floods a target with traffic from multiple sources, causing service disruption. It does not intercept or alter communications between a server and client systems.

q_analyz_netattacks_on-path_02_secp8

A threat actor has positioned themselves in the communication path between two parties in a network. The threat actor is intercepting and possibly altering the communication between the two parties without their knowledge.

Which type of attack does this scenario represent?

Answers:

- Man-in-the-middle attack
- ***On-path attack**
- Distributed denial-of-service (DDoS) attack
- Direct attack

Explanation:

On-path attack is the correct answer. An on-path attack involves a threat actor positioning themselves in the communication path between two parties and intercepting and possibly altering the communication without their knowledge. This is exactly what is described in the scenario.

While a man-in-the-middle attack does involve intercepting and possibly altering communication between two parties, it typically involves the threat actor actively impersonating one or both parties to deceive the other. The scenario does not mention any impersonation, only interception and possible alteration of communication.

Distributed denial-of-service (DDoS) Attack is incorrect because a DDoS attack involves overwhelming a network or service with traffic from multiple sources to cause a denial of service. The scenario does not describe any overwhelming network resources or denial of service.

Direct attack is incorrect because a direct attack involves the threat actor directly attacking the victim's network without the use of third-party servers or specific application protocols. The scenario describes the

q_analyz_netattacks_reflected_secp8

A threat actor has launched an attack against a company's network. The threat actor spoofs the victim's IP address and attempts to open connections with multiple third-party servers. Those servers direct their responses to the victim host, rapidly consuming the victim's available bandwidth.

Which type of attack does this scenario represent?

Answers:

- Distributed denial-of-service (DDoS) attack
- ***Reflected attack**
- Amplified attack
- Direct attack

Explanation:

Reflected attack is the correct answer. In a reflected attack, the threat actor spoofs the victim's IP address and attempts to open connections with multiple third-party servers. Those servers direct their responses to the victim host, which can rapidly consume the victim's available bandwidth.

Distributed denial-of-service (DDoS) attack is incorrect because while a DDoS attack involves multiple hosts launching an attack simultaneously, it does not specifically involve the spoofing of the victim's IP address and the use of third-party servers to direct responses to the victim host.

Amplified attack is incorrect because an amplification attack is a type of reflected attack that targets weaknesses in specific application protocols to make the attack more effective at consuming target bandwidth. The scenario does not mention any specific application protocols being targeted.

Direct attack is incorrect because a direct attack involves the threat actor directly attacking the victim's network without the use of third-party servers. The scenario describes the use of third-party servers to direct responses to the victim host, which is characteristic of a reflected attack.

q_analyz_netattacks_spoof_secp8

A router on the border of your network detects a packet with a source address that is from an internal client, but the packet was received on the internet-facing interface.

This is an example of which form of attack?

Answers:

- Snooping
- Sniffing
- ***Spoofing**
- Spamming

Explanation:

This is an example of spoofing. Spoofing is the act of changing or falsifying information in order to mislead or re-direct traffic. In this scenario, a packet received on the inbound interface cannot receive a valid packet with a stated source that is from the internal network.

Snooping is the act of spying on private information or communications. One type of snooping is sniffing.

Sniffing is the act of capturing network packets in order to examine the contents of communications.

Spamming is sending a victim unwanted and unrequested email messages.

poison_arp_01

What is the MAC address of the first responding device?

Answers:

- 00:00:1B:11:22:33
- 00:00:1b:11:22:33
- 00:00:1b_11:22:33
- 00:00:1B:11:22:33
- 00:00:1B_11:22:33

poison_arp_02

What is the MAC address of the duplicate responding device?

Answers:

- 00:00:1B:33:22:11
- 00:00:1B:33:22:11
- 00:00:1b:33:22:11
- 00:00:1b_33:22:11
- 00:00:1B_33:22:11

q_l_ddos_eh1_01

What indicates that this is a distributed denial-of-service (DDoS) attack?

Answers:

- There are multiple source addresses for the SYN packets with the destination address 128.28.1.1.
- The hex value of the SYN flag is 0x002.
- ***There is a flood of SYN packets without matching SYN-ACK packets.**
- The destination address for all SYN packets is 128.28.1.1.

q_mal_code_bash_sec8

Which of the following statements about Bash is true?

Answers:

- Bash is a command shell and scripting language used only in Windows operating systems.
- Bash was released in 2000 and is rarely used today.
- ***Bash works in the background to execute commands using environment variables.**
- Bash cannot be used to design malware that attacks systems running on Linux's Apache platform.

Explanation:

Bash does work in the background to execute commands using environment variables.

Bash is not used in Windows operating systems. It is used in most Linux distributions and MacOS versions prior to Catalina.

Bash was released in 1989, not 2000, and is still heavily used today.

Malware can actually be designed in Bash to attack systems running on Linux's Apache platform.

q_mal_code_credential_dumping_sec8

You are a cybersecurity analyst at a tech company. You have noticed some unusual activity on a host within your network. You suspect that the host may be compromised and running malicious code.

Which of the following indicators would MOST likely suggest that the host is compromised?

Answers:

- The host is showing signs of shellcode execution.
- The host is showing signs of pivoting/lateral movement.
- ***The host is showing signs of credential dumping.**
- The host is showing signs of persistence mechanisms.

Explanation:

Credential dumping is the correct answer and is a clear sign of a compromised host. Malware often tries to access the credentials file or sniff credentials held in memory by the system process. This is a strong indicator of malicious activity and is more definitive than the other options.

Shellcode execution is a minimal program designed to exploit a vulnerability in the OS or a legitimate app to gain privileges or to drop a backdoor on the host if run as a Trojan. While this could be a sign of a compromised host, it is not as definitive as credential dumping.

Pivoting/lateral movement refers to the procedure used by an attacker to move around within a network, often seeking data assets or trying to widen access by changing the system security configuration. While this could indicate a compromised host, it is not as definitive as credential dumping.

Persistence mechanisms allow the threat actor's backdoor to restart if the host reboots or the user logs off. While this could indicate a compromised host, it is not as definitive as credential dumping.

q_mal_code_powershell_01_sec8

As a system administrator, you notice unusual network activity on a company server. Upon investigation, you discover that a PowerShell script is running in the background.

What type of malware is MOST likely responsible for this activity?

Answers:

- ***Fileless malware**
- Macro virus
- Worm
- Trojan horse

Explanation:

Fileless malware is the correct answer. Fileless malware operates in the memory of the system and does not require a file to run. It is known for using legitimate tools like PowerShell to execute malicious activities. In this scenario, the PowerShell script running in the background is a typical characteristic of fileless malware.

A macro virus is a type of malware embedded in a document and executed when the document is opened. It does not typically use PowerShell for its operations, making this an incorrect answer.

A worm is a type of malware that can replicate itself and spread to other systems, often exploiting system vulnerabilities. While a worm could potentially use PowerShell, the scenario does not provide evidence of replication or spreading, making this an unlikely answer.

A Trojan horse is a type of malware that disguises itself as a normal file or program to trick users into downloading and installing it. While a trojan could potentially use PowerShell, the scenario does not provide evidence of a disguised file or program, making this an unlikely answer.

q_mal_code_powershell_02_sec8

Which of the following statements about PowerShell is true?

Answers:

- PowerShell is a management framework developed by Apple to replace Terminal.
- PowerShell is built on the Java framework and can only be run on Linux operating systems.
- ***PowerShell uses cmdlets to execute commands.**
- PowerShell scripts cannot run in the memory of the system and always need an executable to run.

Explanation:

PowerShell uses cmdlets, which are small scripts that perform certain functions, to execute commands.

PowerShell is a management framework developed by Microsoft, not Apple, and it was designed to replace Command Prompt, not Terminal.

PowerShell is built on the .NET framework, not the Java framework, and it can be run on multiple operating systems, including Windows, MacOS, and Linux.

PowerShell scripts can run in the memory of the system, which means they don't always need an executable to run. This type of malware is known as fileless malware.

q_mal_code_python_sec8

Why is Python a popular programming language for creating Remote Access Trojans (RATs)?

Answers:

- Python is only compatible with Windows operating systems.
- Python's syntax is difficult to learn and understand.
- ***Python can implement libraries that allow RATs to perform various functions.**
- Python scripts are small in size and do not require Python to be installed on the system to run.

Explanation:

Python's ability to implement libraries allows RATs to perform a variety of functions, such as taking screenshots, enabling the webcam and viewing it remotely, making web requests, and making phone calls, makes it a popular programming language for creating RATs.

Python is a versatile language that can be used on most operating systems, including Windows, MacOS, and Linux, not just Windows.

Python is known for its easy-to-learn and readable syntax, which makes it a popular choice for many types of programming, including creating RATs.

One of the main drawbacks to using Python for malware is the file size. Python files are larger than other common languages. Also, Python must be installed on a system for a Python script to run.

q_mal_code_vba_sec8

How can Visual Basic for Applications (VBA) be used to perform malicious attacks?

Answers:

- ***VBA can be used to create a macro virus that opens a shell on the Windows operating system.**
- VBA can be used to disable all security features on a computer system.
- VBA can be used to physically damage the hardware components of a computer.
- VBA can be used to delete all files on a computer system automatically.

Explanation:

VBA can be used to create a macro virus that opens a shell on the Windows operating system. This shell can then be used to perform malicious attacks.

While VBA can be used to perform malicious activities, it cannot directly disable all security features on a computer system.

VBA is a programming language and cannot physically damage the hardware components of a computer.

While a malicious VBA script could potentially delete files, it is not accurate to say that VBA can be used to delete all files on a computer system automatically.

6.6 Analyzing Password Attacks

As you study this section, answer the following questions:

- Where would an attacker gather information to guess a user's password?
- What social engineering technique involves looking through trash?

- What does password spraying help the attacker avoid?
- What is the best defense against password-cracking attempts?
- What is the difference between an online and an offline password attack?

In this section, you will learn to:

- Crack a password using rainbow tables
- Crack a password with John the Ripper

The key terms for this section include:

Term	Definition
Online password attack	An attack where the threat actor interacts with the authentication service directly—a web login form or VPN gateway, for instance.
Offline attack	An attack where once the attacker has obtained a password database, the cracker does not interact with the authentication system.
Social engineering	An activity where the goal is to use deception and trickery to convince unsuspecting users to provide sensitive data or to violate security guidelines.
Shoulder surfing	An eavesdropping technique where the listener obtains passwords or other confidential information by looking over the shoulder of the target.
Brute force attack	A type of password attack where an attacker uses an application to exhaustively try every possible alphanumeric combination to crack encrypted passwords.
Password spraying	A brute force attack in which multiple user accounts are tested with a dictionary of common passwords.
Dictionary attack	A type of password attack that compares encrypted passwords against a predetermined list of possible password values.
Rainbow attack	Similar to dictionary attacks, however, a rainbow attack uses special tables called rainbow tables that have common passwords and the generated hash of each password.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions. • Hashing

	• Salting 2.2 Explain common threat vectors and attack surfaces. • Default credentials • Human vectors/social engineering 2.4 Given a scenario, analyze potential indicators of malicious activity. Password attacks • Spraying • Dictionary • Brute force 5.6 Given a scenario, implement security awareness practices. • User guidance and training ○ Password management
TestOut Security Pro	5.2 Assessment Techniques Analyze password attacks

6.6.1 Password Attacks (Lesson Video)

Transcript:

Passwords are oftentimes the main defense against unauthorized access to a computer system or sensitive data, and attackers have a variety of methods they can use to crack users' passwords. Some of the more common methods are social engineering, brute force attacks, and rainbow attacks. In this lesson, I'll go over how these attacks work and how you can defend against them.

Social engineering is the art of manipulating the weakest link in network security, which is the human element. A skilled attacker can use social engineering to learn a user's password through password guessing, user manipulation, or shoulder surfing.
Guessing a user's password can be difficult, but there are techniques that an attacker can use to make this simpler. The first thing an attacker might do is try default logins and common passwords like password123. An attacker might also use social media to gather personal information such as pet names, family, first car, birthday, and more. Then the attacker uses this information to answer security questions and reset the password itself.
An attacker might also try to interact with a user to trick them into revealing their password. For example, they might call the user pretending to be a technician with an urgent problem for which they need the user's password to resolve. Attackers continue to use this method because, surprisingly, it's still very easy.
An attacker might also use social engineering to gain physical access to an office building. Once inside, they look around for login information that users have written down. Unfortunately, many users write their login information on a sticky pad and put it on their monitor or under their mouse pad. This makes it fairly easy for an attacker to find.
Attackers also go dumpster diving to find these notes that might've been inadvertently thrown away. These people aren't above searching through your trash to gain access to your personal information.
Shoulder surfing is a passive technique where the attacker simply watches a user's keystrokes as they type in their password. This requires access to the user, but it can be accomplished with social engineering.

Your best defense against these kinds of password attacks is education. You must be trained on how to handle these situations. You should know emphatically that no one will ever need to ask for your passwords, and that you should never carelessly write your login information down. With proper user training, attackers have to resort to more technical means to crack passwords.
A brute force password attack can be a very time-consuming proposition. The attacker attempts to guess passwords by using a cracking tool that submits every possible letter, number, and symbol combination in a short amount of time. These attacks can be carried out online or offline.
An online brute force attack requires the attacker to submit passwords with the same user interface while the target is up and running. If the attacker targets a website, they keep submitting logins to the site interface. If the attacker targets a computer, they submit the logins to the login screen.
The best defense against online brute force attacks is to implement lockout policies. This means that if someone enters an incorrect password multiple times in a short period, the system locks them out for a certain amount of time.
An offline brute force attack requires the attacker to somehow steal the password file and then run their attacks on it. This potentially gives the attacker unlimited attempts to run their attack, but it's much more difficult to pull off because they have to first get that password file.
Password spraying is another type of brute force attack that someone can use to get around lockout policies. Instead of attacking one account, the attacker submits a password to multiple user accounts. If none of those work, the attacker submits a different password to multiple accounts. Because there's a delay between submitting passwords on each account, the lockout policy is essentially circumvented.
A dictionary attack is yet another type of brute force attack. In this attack, the hacker compiles a huge list of words and phrases that can be used to guess the password. Weak and common passwords, such as password123, are extremely vulnerable to dictionary attacks.
Brute force attacks require a lot of time and computing power to be effective. This is why hackers use programs like John the Ripper, hashcat, Medusa, and Cain & Abel to facilitate the process.
Implementing proper password protocols is your best defense in these scenarios. A strong password should be at least 8 characters, but more is always better. It should also contain upper and lowercase letters, numbers, and symbols. It shouldn't contain common words and phrases. The longer and more complex your password is, the more time and effort it takes to brute-force it. Also, strongly consider using passphrases instead of passwords.
When you store a plaintext password, it's encrypted and generates a hash. Rainbow attacks are similar to dictionary attacks, but instead of trying to match up words and phrases, a rainbow attack uses tables that are already filled with word lists and their hashes. Rainbow attacks require less computing power and are much faster than brute force and dictionary attacks.
A rainbow table is a table of passwords and generated hashes. The hacker uses this table to match the password's hash instead of the password itself.
For example, let's say that a user's password is TestOut, and this is the generated hash. In the hacker's rainbow table, the password SecurityPro generates the same hash value. This is a hash collision, and the hacker could gain access to the system by sending the hash value without having the actual password itself.
The biggest drawback to rainbow tables is their size. Since different systems use different hashing algorithms, a table needs to be created for each target system. The necessary character sets also greatly increase the tables' sizes. A single table can range from 30 gigabytes to over 300.
The best way to defend against rainbow attacks is to salt your hashes. Salting the hash just means that some random characters are added at the beginning or end of the password, which ends up generating a completely different hash. The login server knows which part of the hash is salted, but anyone who intercepts the hash has no idea.
People use programs like rtgen, winrtgen, and RainbowCrack to create rainbow tables and tools like ophcrack to crack passwords using these tables.
That'll wrap up this lesson. In this lesson, we looked at different password attacks you might encounter. We first looked at social engineering methods, which target the users themselves. Then we covered brute force and dictionary attacks. These methods are essentially attempts at guessing passwords. Remember that these methods take quite a long time and usually require sophisticated computing power and knowledge to pull off. Finally, we looked at rainbow attacks. These attacks attempt to match password hashes instead of having to guess the actual password.

6.6.2 Password Attack Facts

This lesson covers the following topics:

- Password attacks

- Social engineering
- Brute force attacks
- Rainbow attacks
- Cracking passwords using rainbow tables

Password Attacks

Passwords are often the main defense against unauthorized access to computer systems and sensitive data. This makes passwords a prime target for attackers. A variety of attack methods have been developed to retrieve passwords. When a user chooses a password, the plaintext value is converted to a cryptographic hash. This means that, in theory, no one except the user (not even the systems administrator) knows the password because the plaintext should not be recoverable from the hash. A password attack aims to exploit the weaknesses inherent in password selection and management to recover the plaintext and use it to compromise an account.

Online Attacks

An online password attack is where the threat actor interacts with the authentication service directly—a web login form or VPN gateway, for instance. An online password attack can show up in audit logs as repeatedly failed logins and then a successful login or as successful login attempts at unusual times or locations. Apart from ensuring the use of strong passwords by users, online password attacks can be mitigated by restricting the number or rate of login attempts and shunning login attempts from known bad IP addresses.

Note that restricting logins can be turned into a vulnerability as it exposes the account to denial of service attacks. The attacker keeps trying to authenticate, locking out valid users.

Offline Attacks

An offline attack means that the attacker has managed to obtain a database of password hashes, such as %SystemRoot%\System32\config\SAM , %SystemRoot%\NTDS\NTDS.DIT (the Active Directory credential store) or /etc/shadow . Once the password database has been obtained, the cracker does not interact with the authentication system. The only indicator of this type of attack (other than misuse of the account in the event of a successful attack) is a file system audit log that records the malicious account accessing one of these files. Threat actors can also read credentials from host memory, in which case the only reliable indicator might be the presence of attack tools on a host.

If the attacker cannot obtain a database of passwords, a packet sniffer might be used to obtain the client response to a server challenge in an authentication protocol. Some protocols send the hash directly; others use the hash to derive an encryption key. Weaknesses in protocols using derived keys can allow for the extraction of the hash for cracking.

Social Engineering

Social engineering is the art of manipulation. In most networks, the weakest link is the human element. Hackers can take advantage of this to gain access to sensitive information, including passwords.

The following table explains some social engineering techniques to be aware of and protect against.

Social Engineering Technique	Description

Password guessing	Password guessing is usually not a very efficient method to crack a password. An attacker may first attempt to use default login information, such as admin/admin, or simple passwords like password123.
	
	If these do not work, the attacker can use publicly available information, such as on a target's social media, to make the process easier. Information such as the following can be used to guess a password or answer security questions and reset a user's password:
	
	• Birthday
	• First car
	• Family information
	  ○ Spouse's name
	  ○ Child's name
	  ○ Important dates
	• Important locations

User manipulation	A common social engineering technique is user manipulation. This involves the attacker interacting with the user to trick the user into revealing the username and password. For example, the attacker may call the target pretending to be from tech support with an urgent problem. The attacker asks for the target's login information to remote in to resolve the issue. User manipulation is a very successful technique and is still used quite often. User training is the best prevention method.
Physical access	An attacker can use social engineering to gain physical access to an office building. Once inside, the attacker can look around for login information that users have written down. Many users have a tendency to write login information on sticky notes and stick the notes on the monitor or place them under the mouse pad.
Dumpster diving	An attacker may dumpster dive (go through the trash) to find important documents or information that has been thrown out. Many users will throw out papers without realizing the importance of the information. Documents should always be shredded to prevent data loss due to dumpster diving.
Shoulder surfing	Shoulder surfing is an eavesdropping technique in which the attacker obtains passwords or other confidential information by looking over the shoulder of a user typing a password.

User education is the best defense against any form of social engineering. Users should be trained so that no one will ever ask for their login information and always be aware of their surroundings.

Brute Force Attacks

In a brute force attack, the attacker attempts to guess the password by using a cracking tool that submits every possible letter, number, and symbol combination in a short amount of time. A brute force password attack can be a very time-consuming attack.

The following table describes some of the brute force attack methods.

Brute Force Attack Method	Description
Online attack	An online brute force attack requires the attacker to submit the passwords using the same user login interface while the target is up and running. For example: • An attacker targeting a website will submit login attempts to the site interface. • An attacker targeting a computer will submit login attempts to the login screen. The best defense against this method is to implement lock-out policies. This means if the incorrect password is entered multiple times in a short period of time, the account will be locked for a specified amount of time.
Offline attack	Offline attacks require the attacker to somehow steal the password file. The attacker can then run attacks against that file with no limitations, such as lock-out policies. This is the ideal method for the attacker but is more difficult because it requires the attacker to steal the password file somehow.
Password spraying	Password spraying is another method that allows the attacker to avoid lockout policies. • Instead of attempting multiple logins using a single user account and different passwords, the attacker will use the same password with multiple user accounts. • The attacker will continue cycling through the user accounts, submitting passwords until a match is found. • Because there is a delay between submitting a password on each account, the lock-out policy can be avoided.
Dictionary attack	In a dictionary attack, the hacker uses a list of words and phrases to try to guess the password. • Dictionary attacks work well if weak passwords are used. • Using longer and uncommon passphrases is the best way to secure data against these attacks.

Some common password-cracking tools that can be used to carry out brute-force attacks are:

- John the Ripper
- Hashcat
- Medusa
- Cain and Abel

Implementing proper password protocols is the best defense against password-cracking attempts. A strong password should:

- Be at least eight characters; more is better.
- Contain upper and lower case letters.
- Contain numbers.

- Contain symbols.
- Not use common words or phrases.

A passphrase is the best option to use instead of a password.

Rainbow Attacks

When a plaintext password is stored, it is encrypted, and a hash is generated.

Rainbow attacks are similar to dictionary attacks, but instead of trying to match the words and phrases, a rainbow attack uses special tables called *rainbow tables* that are already filled with common passwords and their generated hashes. The attacker uses this table to match the hashes instead of the password. Rainbow attacks require less computing power and are much faster than brute-force attacks.

Storing rainbow tables requires a lot of storage. A single rainbow table can range anywhere from 30 GB to over 300 GB. The character set (lower and/or upper case letters, numbers, symbols) will greatly increase the size. A different rainbow table needs to be generated for each encryption algorithm.

The best defense against rainbow attacks is *salting* the hashes. Salting the hash means that random characters are added at the beginning or end of the password. This generates a completely different hash. The login server is programmed to identify the part of the hash that is salted, but anyone intercepting the hash will have no idea, so the hash cannot be decrypted.

Cracking Passwords Using Rainbow Tables

An encrypted plaintext password stored in a hash file can be cracked using rainbow tables. There are several types of programs that can be used to create and crack these types of passwords, such as:

- Rtgen
- Winrtgen
- RainbowCrack
- Ophcrack

As an example, the following table lists a few examples of the commands needed to create and sort a rainbow crack table:

Command	Description
rtgen	This command generates a rainbow table based on the parameters specified by the user. The parameters are: rtgen ***hash_algorithm charset plaintext_len_min plaintext_len_max table_index chain_len chain_num part_index*** Example: **rtgen md5 ascii-32-95 1 7 0 1000 1000 0** • *hash_algorithm* - A hashing algorithm is a mathematical algorithm that can convert an input data array of a certain type and arbitrary length to an output bit string of a fixed length. A rainbow table must be generated for the type of hash algorithm used. Although there are many hash algorithms that can be used, some of the more common are ntlm, md5, and sha1. • *charset* - A charset specifies all the possible characters for the plaintext. Some of the possible charset that can be used include:

	○ Numeric = [0123456789] ○ alpha = [ABCDEFGHIJKLMNOPQRSTUVWXYZ] ○ alpha-numeric = [ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789] ○ loweralpha = [abcdefghijklmnopqrstuvwxyz] ○ loweralpha-numeric = [abcdefghijklmnopqrstuvwxyz0123456789] ○ ascii-32-95 = ascii-32-95 = [!"#$%&'()*+,-./0123456789:;<=>?@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_'abcdefghijklmngpqrstuvwxyz{|}~] • *plaintext_len_min* and *plaintext_len_max* - These two values, such as 1 7, specify the length of the plaintext. The next four parameters are advanced values and are beyond the scope of this lesson. Therefore, only a brief explanation is given here: • *table_index* - Specifies the reduction function. Examples are: 0, 1, 2, 3, 4. Zero is often used as the default. • *chain_len* - This specifies the rainbow chain length. • *chain_num* - This specifies the number of rainbow chains to generate. • *part_index* - The number of files used to store the rainbow table. If a value greater than zero is used, the rainbow table is saved in the number of smaller files specified by the value. As shown in the example above, common values for these four parameters are: **0 1000 1000 0**
rtsort	A rainbow table is an array of rainbow chains. Each rainbow chain has a starting point and an endpoint. The rtsort program sorts the rainbow chains by endpoint to make binary search possible. To sort a rainbow table, use the following command (the period at the end is part of the command): **rtsort .**

After the rainbow table has been created, you are now ready to crack the passwords. This can be done using the **rcrack** command.

The **rcrack** syntax is: **rcrack *path parameter***

The following table lists a few examples of how the **rcrack** command can be used:

Command	Description
rcrack . -l /root/hashes.txt	The **-l** parameter loads the hashes from a file, and each hash is shown on its own line. The hash is shown, followed by the cracked password. Example output: plaintext of 590cb9bZaC590/5b9b4b0/152d2321117 P@ssw0rd plaintext of 400238780e6c41f8f790161e6ed4aafc21 Test_Out@11_Last plaintext of 89BF04763BF91C9EE2DDBE23D735C73OBDD41FF2 NeverLAnd5

rcrack . -h ***hash_value***	The **-h** parameter loads and displays the results for a single hash. Example command: **rcrack . -h 590cb9bZaC590/5b9b4b0/152d2321117** Example output: **590cb9bZaC590/5b9b4b0/152d232111a P@ssw0rd hex:444387**

6.6.3 Using Rainbow Tables (Demo Video)

Transcript:

Nearly every computer system that requires password authentication must contain a password database. Because the information in databases is vulnerable to theft, storing passwords in plain text isn't a very good idea. Therefore, passwords are usually stored with a cryptographic hash.

When they want to crack a password, hackers often use rainbow tables to speed up the process. So, what is a rainbow table, exactly? A rainbow table is basically a table of passwords and the computed matching hashes. The advantage of using rainbow tables is that they save computation time compared to typical brute forcing. The disadvantage is that rainbow tables can take up a lot of storage space. Here, we can look at the website called RainbowCrack, and you can see that some of these are quite large, over half a terabyte in size.
In this demo, we're going to generate our own rainbow table using the RainbowCrack collection. We're going to use Kail Linux and RainbowCrack, which comes pre-installed.
To open the program, we just type 'rtgen', and it starts. Down here, we have a couple of examples. In the example, the first part is the hashing algorithm. Here, they have MD5. Next, we need to specify the character set. They've specified loweralpha numeric, which is the letters A to Z and numbers 0 to 9. 1 and 7 are the length of the passwords, so they're going to be between one and seven characters long. The rest of this information is some advanced configuration settings that aren't in the scope of what we're covering in this demo. You can find more information about the advanced features on the RainbowCrack website.
I'm going to go ahead and create our table. I'll do that by typing 'rtgen md5' numeric for our character set and 1 through 9 for our password length along with the rest of the parameters. Let's press Enter. When I do that, my rainbow table is created. Now, if I want to see where this was located, I can type in 'cd /usr/share/rainbowcrack' and press Enter. From here, I'll type in 'ls', and you can see my rainbow table file.
I'll clear the screen.
Now we have to sort our rainbow table. To do that, I'll type in 'rtsort .' and press Enter. Each rainbow chain has a start point and an end point. The rtsort program sorts the rainbow chains by end point to make binary searches possible. It's very important that you don't interrupt the sorting process.
All right. That's done, and now we're ready to crack a hash with our new table. I'll clear the screen again.
To crack a hash, I'll type in './rcrack .-h'. The dot is to specify our rainbow tables, and the -h is for hash. Now, we need to put in our hash. To do this, I'm going to go to this website, md5 hash generator. I can generate a hash by entering '123' and clicking Generate. I'll copy this hash, minimize this window, go back to my terminal, paste, and press Enter.
My results pop up very quickly. Here, I've found one of one plaintext. It only took a fraction of a second to run. And down here, I have the hash that I typed in. You can see the numbers I used, 123. So, we successfully cracked this hash with our rainbow table.
That's it for this demo. In this demo, we used the command line version of RainbowCrack to create a rainbow table. Then we sorted the table and used an online hash generator to create a hash that we were able to crack with the rainbow table we created.

6.6.4 Crack Password with Rainbow Tables (Simulation)

Scenario

A recent breach of a popular third-party service has exposed a password database. The security team is evaluating the risk of the exposed passwords for the company. The password hashes are saved in the root user's home directory,

/root/captured_hashes.txt . You want to attempt to hack these passwords using a rainbow table. The password requirements for your company are as follows:

- The password must be 12 or more characters in length.
- The password must include at least one uppercase and one lowercase letter.
- The password must have at least one of these special characters: !, ", #, $, %, &, _, ', *, or @.
- All passwords are encrypted using a hash algorithm of either md5 or sha1.

In this lab, your task is to:

- Create md5 and sha1 rainbow tables using **rtgen** .
- Sort the rainbow tables using the **rtsort** command.
- Crack the hashes using the **rcrack** command. You can run **rcrack** on an individual hash or the hash file (/root/captured_hashes.txt).
- Answer the questions.

The type of charset that can be used to create a rainbow table is stored in the **/usr/share/rainbowcrack/charset.txt** file. This file can be viewed using the **cat** command.

Explanation

Complete this lab as follows:

1. Create and sort an md5 and sha1 rainbow crack table.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **rtgen md5 ascii-32-95 1 20 0 1000 1000 0** and press **Enter** to create an md5 rainbow crack table.
    c. Type **rtgen sha1 ascii-32-95 1 20 0 1000 1000 0** and press **Enter** to create a sha1 rainbow crack table.
    d. Type **rtsort .** and press **Enter** to sort the rainbow table.
2. Crack the password hashes using **-l** or **-h** .
    a. To crack the password contained in a hash file, type **rcrack . -l /root/captured_hashes.txt** and press **Enter** .
    This command lists the hashes continued in the hash file and shows the passwords.
    b. To crack the password contained in a hash, type **rcrack . -h *hash_value*** and press **Enter** .
    This command only shows the password for the specified hash.
    c. Repeat step 2b for the remaining hashes.
3. Answer the questions.
    a. From the top right, select **Answer Questions** .
    b. Answer the questions.
    c. Select **Score Lab** .

6.6.5 Crack Passwords (Demo Video)

Transcript:

All penetration testers and ethical hackers need the ability to crack passwords. In this demo, we're going to cover some ways to crack passwords in Linux and Windows.

One of the most popular password cracking programs is John the Ripper. It combines a number of password crackers into one package, autodetects password hash types, and includes a customizable cracker. You can run it against many

encrypted password formats, including several password hash types commonly found in Linux and Windows. You can also run John the Ripper against access passwords for compressed ZIP files and documents.
John the Ripper is a command line tool. There's also a GUI-based version called Johnny, which uses the underlying features of the command line version. Both are included with Kali Linux. We're going to use the command line version.
First, we need to open up a terminal. Now I want to point out the password file. When we do a dictionary attack, our cracking program needs a wordlist (a list of potential passwords) to refer to. So we need to switch to the folder where that password list is stored. By default, on Kali Linux, the wordlist is located in the /usr/share/john folder. To get there, I need to switch to that directory by typing 'cd /usr/share/john' and press Enter. Now, if I type 'ls', I'll be able to see the contents of this folder. I'm looking for this file, the password.lst file.
Now we're going to go back to the root directory. I can do that by just typing in 'cd'. I like to keep my terminal screen cleared off, so I'll 'clear' the screen.
Next, if I type in 'john', I get a list list some of the different options that I can use with John the Ripper. If you need to know something about a command's syntax, this page might help you out.
We're ready to crack our password. You might think that passwords would be stored in a folder called Passwords, but they're not. Linux passwords are kept in the /etc/shadow file. So, to crack the passwords on this system, I just need to type in 'john /etc/shadow' and press Enter. And that only took a few seconds to crack. It says it loaded two different passwords from two different salts. (If you remember, if a hash is salted, that means that there's extra data added to the hashes because there are two identical passwords being used here. Hashes are salted to make passwords unique in the system.)
Down here, we can see that it found the password for the root user.
Depending on your system, this process could take a lot longer. My password just happens to be toward the top of the list in my password.lst file.
Now we're going to jump over to a Windows system and crack a few passwords from there. We're going to look at dictionary attacks, which use a dictionary file, and brute force attacks, which use every potential combination of letters, numbers, and special characters. And remember, there are many hybrid attacks that mix these techniques in various ways, but we're just covering the basics here.
Let's start with a dictionary attack. We're going to use Cain, a popular program that's also known as Cain and Abel. When you install Cain and Abel, Windows will give you all types of warning messages. You'll have to disable your Windows Firewall and your antivirus software. Your antivirus will most likely recognize Cain and Abel as malicious software and quarantine the files required for it to run.
Once you launch Cain and Abel, you'll see all your password cracking tools. We're going to go to the Cracker tab, up here. We're interested in the LM & NTLM hashes. Next, let's go up here and click on this plus sign to import hashes from our computer. Here, under Add NT Hashes, we can add hashes from our local system. We can also check this box to include password history hashes so that if users are required to change passwords every 30 or 45 days, previous password hashes are included during the import.
You can also import hashes from a text file. So, if you have access to another system, you could potentially collect hashes from that system, dump them into a text file, copy them to a USB device, and then load them here and attempt to crack the passwords. You can also import from a SAM database by checking this, here.
Right now, we're just going to import hashes from our local system. You can see that it just loaded all the local accounts I have on this system. If you see "empty" under NT passwords, that means that there's no password for that account. If that field is blank, that means that a password is assigned to the account.
We're going to use a dictionary attack to see if we can retrieve the password for the Mrs. Worley account. I'm going to right-click on that account, hover over Dictionary Attack, and go to NTLM Hashes. Now I'm going to go up into this white space, right-click, and select Add to list.
Before I started this demo, I downloaded a list of passwords from the internet. This file is a list of ten thousand common passwords. It's estimated that thirty percent of all passwords are included in this list. If you look down here, you can see that it's going to scan for passwords.
If we check this box, here, it will look for the exact password. So, if my password is cat, that's what it's going to look for. Below that, it will reverse that password--so, instead of c-a-t, it would look for t-a-c.
Double passwords are just a password typed in twice. You can also have it search for all-lowercase passwords, search for all-uppercase passwords, or even have it put in numbers that commonly replace letters in passwords. With all these combinations, we're really going to use a variation of tens of thousands of different passwords.
Let's click on Start. Depending on your computer's processor, memory, and capability, this process can take a very long time. To save time, I actually edited that file and deleted most of the passwords with this scan. You can see that the scan did finish, but it didn't find that password. That's a good thing; it means my password isn't a dictionary word. Since that didn't work, let's exit and try a brute force attack.

We'll right-click on Mrs. Worley's account, hover over Brute-Force Attack, and select NTLM Hashes. Here's a predefined attack. You can see that if I click here, we can add to our options. We only have lowercase letters and numbers. We also have lowercase with numbers and special characters. I have lowercase, uppercase, numbers, and special characters. I'm going to brute force from one to sixteen characters. This system is a virtual machine with a single processor and not much memory, so it's going to take a long time. Let's click Start and see what happens. Down here, on the left, you can see it's going to take years to cycle through every variation and crack this password. We don't have years and years to wait, so let's stop this process and try something else.

To crack the password faster, we need to customize our attack. I'll get rid of these characters. Let's say I was doing some shoulder surfing and happened to see that Mrs. Worley was typing in a password, and I noticed three things: I saw that she typed 'teache' and that she didn't use any letters from u to z. I also saw that her password wasn't very long. So, we're going to brute force with these letters and use 6 or 7 characters. We also know the first several letters of our password. Of course, you probably wouldn't have all this information. But to save time, I've made these adjustments. I'll go ahead and click Start, and it will attempt to crack this password. You can see it's going to take a few minutes to complete, so I'll pause the demo and let the attack run.

Okay, our scan has completed in just under three minutes. You can see that the password it found was the word "teacher." Of course, knowing the first several letters of the password, like we did here, isn't common. I simplified the attack for time's sake. But it's often possible to gather enough information to brute force a password in a few days or even hours.

Okay, that's it for this demo. In this demo, we introduced John the Ripper and used it to crack a Linux password. Then we switched over to a Windows system and used Cain and Abel to crack a Windows-based password. You can see how, if users leak a little information or use easy passwords, hacking tools can get you into restricted systems and files in no time.

6.6.6 Crack Password Protected Files (Demo Video)

Transcript:

Password-protected files are extremely common, but they're not 100 percent secure. Some files natively support password protection, like PDFs. Other types of files are easy to store in a password-protected zip archive. Today, we're going to extract the password hash from these kinds of files and use John the Ripper to crack the hash. As a reminder, John the Ripper only works with hashes. It doesn't natively know how to work directly with password-protected files, so we need to extract the hash from the file we're interested in.

The process of extracting a hash is similar for most files. The toolkit provided with John the Ripper provides many scripts, but some of them aren't included on Kali Linux by default. We can obtain these tools by cloning the project from GitHub, but I've already downloaded the tools we need. Let's extract the hashes from the zip file we're interested in.

I'll open a terminal and go to the Documents directory. I'll list the folder contents to see what we have in there. I have a file called protected.zip. This is the zip file we'll attempt to crack the password on. We also have a wordlist that we'll use called words.txt.

Most tools that exist to help us extract hashes from files follow the standard for file type, in this case, zip followed by the number 2 and the word "john". So, I'll type, 'zip2john'. This makes it easier to remember what command we need. These tools output the hash straight to the screen, so we use the angle bracket to redirect that output to a file. We will finish our command with 'protected.zip > ziphash'.

Now that we have the hash for the zip file, we can crack it with John the Ripper. We need to specify the format of the hash, the file name, and the wordlist we want to facilitate cracking the hash. So, I'll type, 'john - -format=pkzip ziphash - -wordlist=words.txt'. We have the password. It's a variation of the word password. Now, let's extract the zipped file with that password!

I'll do that from the command line by typing 'unzip protected.zip'. It prompts for the password, so I'll type in the password that we just cracked. I get a message that the file is inflating. Let's list the folder contents. Now I see another file called password.pdf.

This is a PDF file that's also password-protected, so whoever set this up was really concerned about this document's security. If this PDF's password is just as weak as the first one, we can crack it, too. Let's go ahead and do that now. Getting the hash from the PDF is pretty much the same process as getting the hash from the zipped file. Let's type 'pdf2john password.pdf > pdfhash' and press Enter.

Now let's crack the password, just like we did for the zipped file password, by typing 'john - -format=pdf pdfhash - -wordlist=words.txt' and pressing Enter. There's the password! It's the word 'secret". Let's try it to make sure it works. I'll go to my Documents folder and open the PDF. It prompts me for a password, so I'll type in 'secret', and my file opens up.
And that's all there is to it. In this demo, we talked about extracting hashes from password-protected files. We extracted hashes from a zip file and then a PDF file. Then we used John the Ripper to crack those hashes.

6.6.7 Crack a Password with John the Ripper (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You've received a zip file that contains sensitive password-protected files. You need to access these files. The zip file is located in the home directory.

In this lab, your task is to use John the Ripper to:

- Crack the root password on the Linux computer named Support.
- Crack the password of the **protected.zip** file located in the home directory on IT-Laptop.

After John the Ripper cracks the password, it will not crack it again. The results are stored in the john.pot file.

Explanation

Complete this lab as follows:

1. View the current John the Ripper password file.
    a. From the Favorites bar, select **Terminal** .
    b. At the prompt, type **cd /usr/share/john** and press **Enter** .
    c. Type **ls** and press **Enter** .
    d. Type **cat password.lst** and press **Enter** to view the password list.
    e. Type **cd** and press **Enter** to go back to the root.
2. Crack the root password on the Support computer.
    a. Type **john /etc/shadow** and press **Enter** . The password is shown. Can you find it?
    b. Type **john /etc/shadow** and press **Enter** to attempt to crack the Linux passwords again. Notice that it does not attempt to crack the password again. The cracked password is already stored in the john.pot file.
    c. Use alternate methods of viewing the previously cracked password.
        - Type **john /etc/shadow --show** and press **Enter** .
        - Type **cd .john**
        - **cat john.pot** and press **Enter** to view the contents of the john.pot file.
    d. From the top right, select **Answer Questions** .
    e. Answer Question 1.
    f. Minimize the Lab Questions dialog.
3. Open a terminal on the IT-Laptop.
    a. From the top left, select **Floor 1 Overview** .
    b. Under IT Administration, select **IT-Laptop** .
    c. From the Favorites bar, select **Terminal** .
4. Export the contents of the protected.zip file to a text file.
    a. At the prompt, type **ls** and press **Enter** .
    Notice the protected.zip file you wish to crack.
    b. Type **zip2john protected.zip > ziphash.txt** and press **Enter** .
    c. Type **cat ziphash.txt** and press **Enter** to confirm that the hashes have been copied.
5. Using the text file, crack the password of the protected.zip file.

a. Type **john --format=pkzip ziphash.txt** and press **Enter** to crack the password. The password is shown. Can you find it?
b. Type **john ziphash.txt --show** and press **Enter** to show the previously cracked password.
c. From the top right, select **Answer Questions** .
d. Answer Question 2.
e. Select **Score Lab** .

6.6.8 Practice Questions (Section Quiz)

q_pwd_attacks_brute_01_secp8

Which of the following is MOST vulnerable to a brute-force attack?

Answers:

- ***Password authentication**
- Biometric authentication
- Two-factor authentication
- Challenge-response token authentication

Explanation:

Password authentication is the most vulnerable to a brute-force attack. The brute-force attack itself may take a considerable amount of time, especially if the attack is against a single user account or online login prompt rather than a localized copy of a security account's database. However, once the attack is complete, the attacker has all they need to log into the secured system.

The following authentication types were designed to avoid brute-force attacks and are much less vulnerable to a brute-force attack than password authentication:

- Biometric authentication is a cybersecurity process that verifies a user's identity using unique biological traits such as fingerprints, voices, retinas, and facial features. Biometric authentication systems store this information in order to verify a user's identity when that user accesses their account.
- Two-factor authentication (2FA) is a security system that requires two separate, distinct forms of identification in order to access something. The first factor is a password, and the second commonly includes a text with a code sent to your smartphone or some type of biometrics.
- Challenge-response authentication is a set of protocols used to protect digital assets and services from unauthorized users, programs, or activities. While challenge-response authentication can be as simple as a password, it is often something more complex or dynamic, such as a randomly generated request.

q_pwd_attacks_brute_02_secp8

You are using a password attack that tests every possible keystroke for every single key in a password until the correct one is found.

Which of the following technical password attacks are you using?

Answers:

- Keylogger
- Pass-the-hash attack
- ***Brute force attack**
- Password sniffing

Explanation:

In a brute force attack, every password is eventually found because the technique is to test every possible keystroke for every single key in a password until the correct one is found.

Keyloggers log or record every keystroke on the computer keyboard to obtain passwords and other important data.

A pass-the-hash attack is a hacking technique where an attacker uses an underlying NTLM or hash of a user's password to gain access to a server without ever using the actual plaintext password.

Password sniffing is a passive way for attackers to gain access to an account. The sniffer collects data that is in transit in a LAN. If access is gained on one system in a LAN, data can be gathered from data being sent from any other system in the network. The sniffer runs in the background, making it undetectable.

q_pwd_attacks_brute_03_secp8

An organization is experiencing an attack where the attackers break into the premises or cabinets by forcing a gateway or locks.

What BEST describes the observed attack?

Answers:

- ***Brute force**
- Downgrade
- Collision
- Birthday

Explanation:

A brute force physical attack can include smashing a hardware device to perform physical denial-of-service (DoS) or breaking into the premises or cabinets by forcing a gateway or locks.

A downgrade attack is a cryptographic attack that forces a server or client to use a lower specification protocol with weaker ciphers and key lengths.

A collision attack is a cryptographic attack where a weak cryptographic hashing function or implementation allows the generation of the same digest value for two different plaintexts.

A birthday attack is a cryptographic attack that exploits weaknesses in the mathematical algorithms used to encrypt passwords. It takes advantage of the probability of different password inputs producing the same encrypted output.

q_pwd_attacks_brute_04_secp8

While on patrol, a security guard discovers signs of forced entry on a window leading into the HVAC control room. Upon closer examination, the security guard notices hardly any lights or sounds of machinery working in the room.

These are indicators of what type of attack?

Answers:

- ***Brute force**

- On-path
- Distributed denial-of-service
- Network

Explanation:

A brute force attack is an attack against a physical aspect of a facility containing the physical hardware. Forcing a window and damaging HVAC equipment are excellent indicators of a brute force attack.

An on-path attack is where a malicious attacker can place themselves between two recipients to intercept messages sent via the network.

A distributed denial-of-service (DDoS) attack consumes a system's resources, so legitimate requests cannot go through.

A network-based attack comes over the network and does not require being physically present at the hardware's location to interfere with it.

q_pwd_attacks_dumpster_secp8

You are cleaning your desk at work. You toss several stacks of paper in the trash, including a sticky note with your password written on it.

Which of the following types of non-technical password attacks have you enabled?

Answers:

- Shoulder surfing
- Social engineering
- ***Dumpster diving**
- Password guessing

Explanation:

Dumpster diving relies on finding sensitive information that has been discarded in garbage cans, dumpsters, or other unsecured places that create access for attackers.

Shoulder surfing is watching and recording a password, pin, or access code that is being entered by someone nearby.

Social engineering relies on human error. It works by feigning trustworthiness to convince someone to give the attacker access.

Password guessing happens when someone is able to easily guess a password, typically because it is very common, like a pet's name or a hobby.

q_pwd_attacks_imperson_secp8

Carl receives a phone call from a woman who states she is calling from his bank. She tells him that someone has tried to access his checking account, and she needs him to confirm his account number and password to discuss further details. He gives her his account number and password.

Which of the following types of non-technical password attack has occurred?

Answers:

- Shoulder surfing
- ***Social engineering**
- Dumpster diving
- Password guessing

Explanation:

Social engineering relies on human error. It works by feigning trustworthiness to convince someone to share information.

Shoulder surfing is watching and recording a password, pin, or access code that is being entered by someone nearby.

Dumpster diving relies on finding sensitive information that has been discarded in garbage cans, dumpsters, or other unsecured places that create access for attackers.

Password guessing happens when someone is able to easily guess a password, typically because it is very common, like a pet's name or a hobby.

q_pwd_attacks_offline_secp8

A hacker successfully exfiltrates a database of user passwords and attempts to gain access to it as the hacker can now go around the authentication system.

What type of attack has the hacker achieved?

Answers:

- Password spraying
- Brute force
- Dictionary
- ***Offline**

Explanation:

An offline attack means the hacker has obtained a database of password hashes. The hacker can then perform attacks against this offline database in an attempt to compromise the encryption.

Password spraying is a horizontal brute force online attack. In password spraying, the hacker chooses one or more common passwords and tries them in conjunction with multiple usernames.

A brute force attack attempts every possible combination in the output space to match a captured hash and derive the plaintext that generated it.

A dictionary attack occurs when there is a good chance of guessing the likely value of the plaintext, such as an uncomplicated password.

q_pwd_attacks_pass_crack_secp8

You want to check a server for user accounts that have weak passwords.

Which tool should you use?

Answers:

- ***John the Ripper**
- Retina
- OVAL
- Nessus

Explanation:

John the Ripper is a password-cracking tool. Password crackers perform cryptographic attacks on passwords. Use a password cracker to identify weak passwords or passwords protected with weak encryption.

Nessus and Retina are vulnerability scanners. While vulnerability scanners check for default user accounts and often check for accounts with blank passwords, they typically do not include password-cracking features to test for weak passwords.

The Open Vulnerability and Assessment Language (OVAL) is an international standard for testing, analyzing, and reporting the security vulnerabilities of a system.

q_pwd_attacks_rainbow_01_secp8

Which of the following password attacks uses preconfigured matrices of hashed dictionary words?

Answers:

- Dictionary attack
- Brute-force attack
- Hybrid attack
- ***Rainbow table attack**

Explanation:

A rainbow table attack applies hashing algorithms to every word in a dictionary (sometimes including hybrids or passwords accumulated in brute force techniques). The algorithm then saves the results in a table or matrix. An encrypted password is compared to the pre-computed hashed passwords in the matrix until a match is found.

A dictionary attack tries known words (such as from a dictionary).

A brute force attack works through all possibilities until the password is cracked.

A hybrid attack adds appendages to known dictionary words (for example, 1password, password07, and p@ssword1).

q_pwd_attacks_rainbow_02_secp8

Which of the following strategies can protect against a rainbow table password attack?

Answers:

- ***Add random bits to the password before hashing takes place.**
- Encrypt the password file with one-way encryption.
- Educate users to resist social engineering attacks.
- Enforce strict password restrictions.

Explanation:

Some authentication protocols send password hashes between systems during the authentication process. Rainbow table attacks apply hashing algorithms to every word in a dictionary (sometimes including hybrids or passwords accumulated in brute force techniques) in an attempt to match hashed passwords. To protect against this type of attack, you can salt the hash by adding random bits to the password before hashing takes place, thereby producing an entirely different hash value for the password. Because the hacker does not know the extra random bits, the rainbow table is of no value.

The password file should be encrypted. Rainbow attacks do not work by accessing the password file but by capturing hashed passwords being transmitted on the network.

Users should be educated about social engineering attacks, but there is no connection between social engineering and rainbow table attacks.

Enforcing strict password restrictions might actually weaken network security if you do not educate users about proper procedures that protect login credentials.

q_pwd_attacks_rainbow_03_secp8

Jack is tasked with testing the password strength for the users of an organization. He has limited time and unlimited storage space.

Which of the following would be the BEST password attack for him to choose?

Answers:

- ***Rainbow attack**
- Dictionary attack
- Brute-force attack
- Keylogger attack

Explanation:

Rainbow attacks are like dictionary attacks, but instead of endlessly testing dictionary lists, they use tables precomputed with word lists and their hashes. This is much quicker than a dictionary attack or a brute-force attack. The biggest drawback to rainbow tables is their size. A single table can range from 30 gigabytes to over 300.

In a dictionary attack, word lists often taken straight from dictionaries are tested against password databases. Besides all the standard words you find in a dictionary, these lists usually include variations on words that are common for passwords, like using the word pa$$word. Lists can also include simple keyboard finger rolls, like q-w-e-r-t1234. The downside to this attack is this process can take a very long time.

In a brute-force attack, every possible keystroke is tested for each single key in a password until the correct one is found. The disadvantages of this type of attack are that it takes a huge amount of processing power to execute and it is very time-consuming.

Keyloggers record every stroke on the computer keyboard. Still, they must either be installed manually on each computer with the hardware option, or every user has to open an email attachment to install the software option. Both processes are very time-consuming.

q_pwd_attacks_salting_secp8

Which of the following techniques involves adding random bits of data to a password before it is stored as a hash?

Answers:

- Password sniffing
- Pass-the-hash attack
- ***Password salting**
- Keylogging

Explanation:

Password salting is adding random bits of data to a password before it is stored as a hash, making password cracking much more difficult.

Password sniffing is a passive way for attackers to gain access to an account. The sniffer collects data that is in transit in a LAN.

A pass-the-hash attack is a hacking technique where an attacker uses an underlying NTML or hash of a user's password to gain access to a server without ever using the actual plaintext password.

Keylogging is recording every stroke on the computer keyboard.

q_pwd_attacks_shoulder_secp8

Which of the following BEST describes shoulder surfing?

Answers:

- Guessing someone's password because it is so common or simple.
- Finding someone's password in the trash can and using it to access their account.
- Giving someone you trust your username and account password.
- ***Someone nearby watching you enter your password on your computer and recording it.**

Explanation:

Shoulder surfing is watching and recording a password, pin, or access code that is being entered by someone nearby.

Password guessing happens when someone is able to easily guess a password, typically because it is very common, like their pet's name or their hobby.

Dumpster diving relies on finding sensitive information that has been discarded in garbage cans, dumpsters, or other unsecured places that create access for attackers.

Social engineering relies on human error. It works by convincing someone to give the attacker access because they trick them into trusting them.

q_pwd_attacks_spraying_secp8

An organization notices an external actor trying to gain access to the company network. The attacker is not targeting a specific account but rather using the same password across a vast range of usernames in hopes that one might be correct.

What type of attack BEST describes this scenario?

Answers:

- ***Spraying**
- Brute force
- Rainbow table
- Dictionary

Explanation:

In a password spraying attack, the attacker tries a small number of commonly used passwords on many different accounts, attempting to bypass account lockouts that would generally affect brute-force attacks.

Brute-force attacks involve an attacker trying many passwords or passphrases with the hope of eventually guessing correctly. Brute-force attacks usually target one account.

Rainbow table attacks use precomputed tables for reversing cryptographic hash functions for cracking password hashes. They do not involve the widespread, low-frequency approach described in the scenario.

Dictionary attacks are typically applied to one account (or a small number of accounts) and are a parent of attacks like password attacks. However, it specifically mentions that the attack occurs across multiple accounts.

q_pwd_attacks_user_manipulation_secp8

Which social engineering technique involves the attacker interacting with the user to trick them into revealing their username and password?

Answers:

- Password guessing
- Physical access
- ***User manipulation**
- Dumpster diving

Explanation:

User manipulation is the correct answer. In user manipulation, the attacker interacts directly with the user, often pretending to be someone they trust, to trick them into revealing their username and password.

Password guessing involves the attacker trying to guess the user's password using default login information or publicly available information about the user. It does not involve direct interaction with the user.

Physical access involves the attacker gaining physical access to an office building or a user's workspace to look for written-down login information. It does not involve tricking the user into revealing their credentials.

Dumpster diving involves the attacker going through the user's trash to find discarded documents or information that might reveal the user's login information. It does not involve direct interaction with the user.

rainbow_table_01

What is the password for hash 202cb962ac59075b964b07152d234b70?

Answers:

- 123

rainbow_table_02

What is the password for hash 400238780e6c41f8f790161e6ed4df3b?

Answers:

- MaryHad_A_Sm@ll_Lamb

rainbow_table_03

What is the password for hash 89BF04763BF91C9EE2DDBE23D7B5C730BDD41FF2?

Answers:

- DisneyL@nd3

rainbow_table_04

How many of the passwords found meet the company's password requirements?

Answers:

- 0
- ***1**
- 2
- 3

john_ripper_question1

What is the password for the root user of the Kali Linux computer?

Answers:

- 1worm4b8

john_ripper_question2

What is the password for the protected.zip file?

Answers:

- p@ssw0rd
- p@sswOrd

7.0 Vulnerability Management

7.1 Vulnerability Management

As you study this section, answer the following questions:

- Why is vulnerability management critical to cybersecurity strategy?
- In what ways do legacy and end-of-life systems increase the risk of vulnerabilities?
- What factors should be considered when conducting a vulnerability analysis?
- Why is reporting crucial for vulnerability management?
- How do threat feeds help with vulnerability management?

In this section, you will learn to:

- Explore end-of-life software/hardware

The key terms for this section include:

Term	Definition
Vulnerability management	Identifying and managing the risks to a network, including the operating system, applications, and other components of an organization's IT operations.
Vulnerability scan	Utilizes automated scanning processes to identify and evaluate potential issues.
Threat feed	Real-time, continuously updated sources of information about potential threats and vulnerabilities.
Penetration testing	We pay you to hack into systems to make sure people can't hack into systems (a very off quote from Sneakers).

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.1 Compare and contrast various types of security controls. • Control types ○ Compensating 2.3 Explain various types of vulnerabilities. • Operating system (OS)-based

- Hardware
    - Firmware
    - End-of-life
    - Legacy

4.3 Explain various activities associated with vulnerability management.

- Identification methods
    - Vulnerability scan
    - Threat feed
        - Open-source intelligence (OSINT)
        - Proprietary/third-party
        - Information-sharing organization
        - Dark web
    - Penetration testing
    - Responsible disclosure program
        - Bug bounty program
    - System/process audit
- Analysis
    - Confirmation
    - Prioritize
    - Common Vulnerability Scoring System (CVSS)
    - Common Vulnerability Enumeration (CVE)
    - Vulnerability classification
    - Exposure factor
    - Environmental variables
    - Industry/organizational impact
    - Risk tolerance
- Vulnerability response and remediation
    - Patching
    - Insurance
    - Segmentation
    - Compensating controls
    - Exceptions and exemptions
- Validation of remediation
    - Rescanning
    - Audit
    - Verification

4.4 Explain security alerting and monitoring concepts and tools.

- Tools
    - Vulnerability scanners

4.8 Explain appropriate incident response activities.

- Threat hunting

4.9 Given a scenario, use data sources to support an investigation.

- Data sources
    - Vulnerability scans

TestOut Security Pro	3.1 Harden computer systems 3.1.4 Configure Windows Update 5.2 Assessment techniques 5.2.3 Scan for vulnerabilities

7.1.1 Vulnerability Management (Lesson Video)

Transcript:

A vulnerability is a weakness that could be triggered accidentally or exploited intentionally to cause a security breach. Efficient vulnerability management plays a vital role in securing your organization's network. In this video, we'll discuss steps that can be taken to manage vulnerabilities effectively.

Let's start by looking at identification. Various sources can be used to identify vulnerabilities.
Vulnerability scans are automated tools that scan your system for known vulnerabilities. These scans comprehensively analyze your system's current security status, flagging any areas that require immediate attention or improvement.
Threat feeds provide real-time updates on cyber threats across the world. These feeds can provide information such as suspicious domains, known malware, known malicious IP addresses, and more. Security defenders can use this information to stay on top of the latest threats that might impact them.
Open-source intelligence gathers data from publicly available sources. At the same time, proprietary third-party feeds offer unique intelligence from private companies. Information-sharing organizations offer sector-specific threat intelligence. Dark web intelligence can also provide insights into underground criminal activities, including selling exploit codes and stolen data.
Penetration testing, often called "pen testing," is a simulated attack that can be used to test your defenses. They can help to detect security loopholes before malicious actors find them. A responsible disclosure program encourages external individuals to report vulnerabilities. Bug bounty programs take responsible disclosure one step further by offering recognition or compensation to individuals who report vulnerabilities. Lastly, system and process audits help to identify deviations from established policies. These audits can help to highlight non-compliance issues and pinpoint inefficiencies and opportunities for improvement in current processes.
Next is threat hunting. This involves proactively searching through networks to detect and isolate advanced threats evading security solutions. It's about being proactive instead of reactive, staying one step ahead instead of waiting for an attack to happen.
Threat hunting involves using automated and manual techniques to detect cyber threats. It begins with a hypothesis about a potential threat, which is then investigated using data analysis, system logs, and security tools. The goal is to detect hidden threats, understand the extent of any breaches, and identify ways to prevent further attacks. This proactive approach requires an in-depth understanding of an organization's network, including all internal and external connections, system configurations, and data flows. It also requires an understanding of current threat trends and tactics.
Automated tools can scan and monitor network traffic, but the human element is crucial in threat hunting. Cybersecurity professionals use their experience and intuition to interpret the data, identify anomalies, and track potential threats.
Once vulnerabilities have been identified, they must be analyzed to assess their potential impact. Vulnerability analysis helps prioritize remediation efforts by identifying the most critical vulnerabilities that pose the most significant risk to an organization.
Prioritization is typically based on factors such as a vulnerability's severity, the ease of exploitation, and the potential impact of an attack. Prioritizing vulnerabilities helps an organization focus limited resources on addressing the most significant threats first.
A standard ranking system is the Common Vulnerability Scoring System, or CVSS. The CVSS determines the vulnerability risk based on three metrics: base, temporal, and environmental. Base metrics describe a vulnerability's unique characteristics. Temporal metrics describe its changeable attributes, and environmental metrics tell what vulnerabilities are only present in specific environments or implementations.
Scan results often list the vulnerabilities by their CVE code. CVE stands for Common Vulnerabilities and Exposures. This is a free and publicly available list of standardized identifiers for known software vulnerabilities and exposures. There are currently 94 CVE numbering authorities from 16 countries, which provides a baseline for evaluation. CVE also provides

standardization, which allows data exchange for cybersecurity automation and aids professionals as they determine the best assessment tools for themselves.
Depending on the vulnerability analysis results, appropriate steps should be taken to mitigate the identified risks. Some standard remediation practices include patching, cybersecurity insurance, segmentation, and compensating controls.
Patching involves updating software or systems to fix security vulnerabilities. The process includes installing 'patches' or updates issued by software vendors to fix identified security loopholes. Cybersecurity insurance is a form of cover that offers a financial safety net in case of cyber-attacks or data breaches. It can cover costs associated with recovery, including incident response, data recovery, legal liability, and customer notification.
Segmentation refers to dividing a network into smaller parts or segments. By doing this, organizations can isolate parts of their network, limiting an attacker's ability to move laterally within the system. Compensating controls are alternative security measures implemented when a standard control isn't feasible. These controls help achieve the same security objective and reduce the risk to an acceptable level if the primary control cannot be implemented.
That's it for this lesson. In this lesson, we discussed vulnerability management. We talked about various methods that can be used for identifying vulnerabilities. We discussed how threat hunting, threat analysis, and threat response and remediation can play a vital role in securing an organization's network.

7.1.2 Vulnerability Type Facts

This lesson covers the following topics:

- Vulnerability management
- Device and operating system vulnerabilities
- Legacy and end-of-life (EOL) systems
- Firmware vulnerabilities

Vulnerability Management

Vulnerability management is critical to any organization's cybersecurity strategy, encompassing identifying, evaluating, treating, and reporting security vulnerabilities in operating systems, applications, and other components of an organization's IT operations. Vulnerability management may involve patching outdated systems, hardening configurations, or upgrading to more secure versions of operating systems. For applications, it might include code reviews, security testing, and updating third-party libraries.

Vulnerability scanning is a crucial component of this process, with specialized tools utilized to identify potential weaknesses in an organization's digital assets automatically. These tools scan for known vulnerabilities such as open ports, insecure software configurations, or outdated versions. Post scanning, analysis is performed to validate, classify, and prioritize the identified vulnerabilities for remediation based on factors such as the potential impact of a breach, the ease of exploiting the vulnerability, and the importance of the asset at risk. This continuous cycle of assessment and improvement helps organizations maintain safe and secure computing environments.

Device and Operating System Vulnerabilities

Operating systems (OS) are one of the most critical components of any infrastructure, so vulnerabilities in an OS can lead to significant problems when successfully exploited. Microsoft Windows has an extensive feature set and broad user base, especially among large organizations and governments. Its vulnerabilities often include buffer overflows, input validation problems, and privilege flaws typically exploited to install malware, steal information, or gain unauthorized access. Windows is an essential target for attackers because of its large install base. Large corporations and governments heavily depend upon it, which compounds the significance of its vulnerabilities.

Apple's macOS vulnerabilities often stem from its UNIX-based architecture, and weaknesses generally appear in access controls, secure boot processes, and third-party software. Apple macOS has a smaller user base than Windows, but its popularity has grown significantly. Generally, macOS is perceived as being 'safer' than other operating systems, which can lead to complacency.

Linux is a prevalent server OS but can also be used as a desktop or mobile OS. The open-source nature of Linux and the large community of active developers support rapid development. This generally results in quick identification and repair of vulnerabilities. Kernel vulnerabilities, misconfigurations, and unpatched systems are common issues in Linux. Despite its reputation for security, its widespread use in the cloud and server infrastructure makes Linux vulnerabilities especially significant.

The widespread adoption of Mobile OS like Android and iOS and their increasing use as primary computing platforms instead of traditional computers make them valuable targets for attack and exploitation. Android is open source, like Linux, resulting in similar benefits and problems. Additionally, Android OS is fragmented among different manufacturers and versions, resulting in inconsistent patching and update support. iOS, while not open source like Android, has also been impacted by several significant vulnerabilities.

The significance of OS vulnerabilities cannot be overstated, especially as specialized embedded systems, such as IoT, are added to our surroundings. Each system runs specialty operating systems and introduces vulnerabilities and potential pathways into corporate infrastructures.

Examples of OS vulnerabilities:

- **Microsoft Windows** — One of the most notorious vulnerabilities in Windows history was the MS08-067 vulnerability in Windows Server Service. This vulnerability allowed remote code execution if a specially crafted packet was sent to a Windows server. This vulnerability was exploited by the Conficker worm in 2008, which infected millions of computers worldwide. Additionally, MS17-010 represents a significant and critical security update released by Microsoft in March 2017. This update addressed multiple vulnerabilities in Microsoft's implementation of the Server Message Block (SMB) protocol (a network file-sharing protocol) that could allow remote code execution (RCE). These vulnerabilities, if exploited, could allow an attacker to install programs, view, change, or delete data, or create new accounts with full user rights.
- The significance of MS17-010 is tied closely to the EternalBlue exploit, which leveraged the vulnerabilities in early versions of the SMB protocol for malicious purposes. The most famous misuse of EternalBlue was during the WannaCry ransomware attack in May 2017, where it was used to propagate ransomware across networks worldwide, leading to massive damage and disruption. This event underlined the critical importance of timely system patching and reinforced the potential global impact of such vulnerabilities.
- **macOS** — In 2014, a significant vulnerability called "Shellshock" affected all Unix-based systems, including macOS. It allowed attackers to potentially gain control over a system due to a flaw in the Bash shell. Though it originated from a component in Unix systems, its impact was felt in macOS due to its Unix-based architecture.
- **Android** — The Stagefright vulnerability discovered in 2015 is a prominent example of Android. It allowed attackers to execute malicious code on an Android device by sending a specially crafted MMS message. This issue was particularly severe due to the ubiquity of the vulnerable component (the Stagefright media library) across Android versions and devices.
- **iOS** — In 2019, Google's Project Zero team discovered a series of vulnerabilities in iOS that nation-state attackers were abusing. These "watering hole" attacks took advantage of several vulnerabilities to gain full access to a device by having the victim visit a malicious website.
- **Linux** — The Heartbleed bug in 2014 was a severe vulnerability in many Linux systems' OpenSSL cryptographic software library. The vulnerability allowed attackers to read the system’s memory running the OpenSSL software's vulnerable versions, compromising the secret keys to protect data.

Additional concerns arise in mobile operating systems due to factors like the diversity of devices and operating system versions, bypassing operating system protections, and using apps downloaded outside official app stores. Identifying and managing vulnerabilities often involves keeping operating systems updated with the latest patches, hardening system configurations, carefully managing user privileges, and controlling software applications.

Legacy and End-of-Life Systems

Hardware vulnerabilities, particularly those associated with end-of-life and legacy systems, present considerable security challenges for many organizations, as patches or fixes for vulnerabilities are either unavailable or difficult to apply. End-of-life (EOL) and legacy systems share a common characteristic: they are both outdated. EOL systems may be legacy systems, and some are also EOL.

The manufacturer or vendor no longer supports EOL systems, so they do not receive updates, including critical security patches. This makes them vulnerable to newly discovered threats. Conversely, while still outdated, the vendor may still fully support legacy systems.

An EOL system is a specific product or version of a product that the manufacturer or vendor has publicly declared as no longer supported. It is also possible for open-source projects to be abandoned by the maintainers. An EOL system can be a hardware device, a software application, or an operating system. Products should be replaced or updated before they reach EOL status to ensure they remain supported by their vendors and receive critical security patches. Notable EOL product examples include the Windows 7 and Server 2008 operating systems, which stopped receiving updates in January 2020. These systems are significantly more vulnerable to attacks due to the absence of security patches for new vulnerabilities. Despite their EOL status, they are still in use in many environments.

Many devices (peripheral devices especially) remain on sale with known severe vulnerabilities in firmware or drivers and no possibility of vendor support to remediate them, especially in secondhand, recertified, or renewed/reconditioned marketplaces. Examples include recertified computer equipment, consumer-grade and recertified networking equipment, and various Internet of Things devices.

Legacy systems typically describe outdated software methods, technology, computer systems, or application programs that continue to be used despite their shortcomings. Legacy systems often remain in use for extended periods because the organization's leadership recognizes that replacing or redesigning them will be expensive or pose significant operational risks from complexity. The term "legacy" does not necessarily mean that the vendor no longer supports the system but rather that it represents hardware and software methods that are no longer popular and often incompatible with newer architectures or methods. Legacy systems often remain in use because they operate with sufficient reliability, have been incorporated into many critical business functions, and are familiar to long-tenured staff.

Assessing the risks associated with using EOL and legacy products, such as lack of updates, lack of support, and compatibility issues with newer systems, is crucial. EOL and legacy product replacements must continue to meet the organization's requirements, maintain compatibility with existing infrastructure, and support reliable data migration. Selection criteria must consider the availability of vendor support, device warranty details, and marketplace performance/reputation. Transitioning costs must be carefully assessed, too, including licensing, hardware upgrades, and professional service implementation fees. The work to transition away from EOL and legacy products must minimize disruptions and ensure long-term sustainability.

Firmware Vulnerabilities

Firmware is the foundational software that controls hardware and can contain significant vulnerabilities. For instance, the Meltdown and Spectre vulnerabilities identified in 2018 impacted almost all computers and mobile devices. The exposure was associated with the processors used inside the computer and allowed malicious programs to steal data as it was being processed. Another vulnerability, "LoJax," discovered in the Unified Extensible Firmware Interface (UEFI) firmware in 2018, enabled attackers to persist on a system even after a complete hard drive replacement or OS reinstallation. End-of-life (EOL) hardware vulnerabilities arise when manufacturers cease providing product updates, parts, or patches to the firmware.

7.1.3 Vulnerability Identification Methods Facts

This lesson covers the following topics:

- Vulnerability scan
- Threat feed
- Other vulnerability identification methods

Vulnerability Scan

Vulnerability scanning is a fundamental task within a vulnerability management program. It utilizes automated scanning processes to identify and evaluate potential issues. Vulnerability scans focus on networks, operating systems, applications, and other functional areas to detect known vulnerabilities, including insecure configurations, outdated software versions, or missing security patches. Regular scanning is required to maintain an accurate picture of an organization's security posture and to identify new vulnerabilities.

Threat feeds play a vital role in enhancing the effectiveness of vulnerability management by providing real-time information about emerging threats and newly discovered vulnerabilities. Threat feeds aggregate data from various sources, including cybersecurity researchers, vendors, and global security communities. These are integrated into vulnerability scanning tools to improve their detection capabilities. By leveraging threat feeds, organizations can stay ahead of the threat landscape, enabling them to prioritize and address the most critical vulnerabilities before attackers can exploit them.

Threat Feed

Another essential element of vulnerability management is the use of threat feeds. These are real-time, continuously updated sources of information about potential threats and vulnerabilities, often gathered from multiple sources. By integrating threat feeds into their vulnerability management practices, organizations can stay aware of the latest risks and respond more swiftly.

Threat feeds are pivotal in vulnerability scanning by providing real-time, continuous data about the latest vulnerabilities, exploits, and threat actors. These feeds serve as a valuable resource for enhancing the organization's threat intelligence and enabling quicker identification and remediation of potential vulnerabilities. They integrate data from various sources, including security vendors, cybersecurity organizations, and open-source intelligence, to comprehensively view the threat landscape.

Common threat feed platforms include AlienVault's Open Threat Exchange (OTX), IBM's X-Force Exchange, and Recorded Future. These platforms gather, analyze, and distribute information about new and emerging threats, providing actionable intelligence that can be incorporated into an organization's vulnerability management practices and sometimes directly into security infrastructure tools to provide up-to-the-minute protections.

Threat feeds significantly improve vulnerability identification by providing timely information and context about new threats that traditional vulnerability scanning does not provide. Threat feeds offer information that helps organizations focus their remediation efforts on the most relevant and potentially damaging vulnerabilities first. This proactive approach can significantly reduce the time between discovering a vulnerability and its remediation, thus minimizing the organization's exposure to potential attacks.

The following table describes several types of threat feeds:

Threat intelligence platforms and feeds are supplied as one of three different commercial models:

Threat Feed Type	Description
Third-party threat feeds	Open-source and proprietary threat feeds provide valuable real-time information on the latest cyber threats and vulnerabilities. Both feed types aggregate data from various sources and can be integrated into an organization's security infrastructure, contributing to a proactive cybersecurity strategy. The choice between open-source and proprietary threat feeds often comes down to essential attributes. Open-source feeds, such as those provided by the Cyber Threat Alliance or the MISP threat-sharing platform, are typically free and accessible to all, making them a cost-effective solution for smaller organizations or those with limited budgets. However, they may lack the depth, breadth, or sophistication of analysis found in proprietary feeds. Proprietary threat feeds often provide more comprehensive information and advanced analytic insights. However, these feeds come at a cost, and the return on investment will depend on an

Threat Feed Type	Description
	organization's specific needs, risk profile, and resources. Some organizations may use a combination of both open-source and proprietary feeds to achieve a balance of cost and coverage. The outputs from the primary research undertaken by threat data feed providers and academics can take three main forms: • Behavioral Threat Research — is narrative commentary describing examples of attacks and TTPs gathered through primary research sources. • Reputational threat intelligence — is lists of IP addresses and domains associated with malicious behavior, plus signatures of known file-based malware. • Threat Data — is computer data that can correlate events observed on a customer's networks and logs with known TTP and threat actor indicators. Threat data can be packaged as feeds integrating with a security information and event management (SIEM) platform. These feeds are usually described as cyber threat intelligence (CTI) data. The data is not a complete security solution; the threat data must be correlated with observed data from customer networks to produce actionable intelligence. This type of analysis is often powered by the SIEM's artificial intelligence (AI) features. • Closed/proprietary — is where threat research and CTI data are available as a paid subscription to a commercial threat intelligence platform. The security solution provider will also make the most valuable research available early to platform subscribers through blogs, white papers, and webinars. Some examples of such platforms include the following: ◦ IBM X-Force Exchange (exchange.xforce.ibmcloud.com) ◦ Mandiant's FireEye (mandiant.com/advantage/threat-intelligence) ◦ Recorded Future (recordedfuture.com/platform/threat-intelligence)
Open-source intelligence	Open-source intelligence (OSINT) describes collecting and analyzing publicly available information and using it to support decision-making. In cybersecurity operations, OSINT is used to identify vulnerabilities and threat information by gathering data from many sources such as blogs, forums, social media platforms, and even the dark web. This can include information about new types of malware, attack strategies used by cybercriminals, and recently discovered software vulnerabilities. Security researchers can use OSINT tools to automatically collect and analyze this information, identifying potential threats or vulnerabilities that could impact their organization. Some standard OSINT tools include Shodan for investigating internet-connected devices, Maltego for visualizing complex networks of information, Recon-ng for web-based reconnaissance activities, and theHarvester for gathering emails, subdomains, hosts, and employee names from different public sources.
Information-sharing organization	Threat feed information-sharing organizations are collaborative groups that exchange data about emerging cybersecurity threats and vulnerabilities. These organizations collect, analyze, and disseminate threat intelligence from various sources, including their members, security researchers, and public sources. Members of these organizations, often composed of businesses, government entities, and academic institutions, can benefit from the shared intelligence by gaining insights into the latest threats they might not have access to individually. They can use this information to fortify their systems and respond swiftly to emerging threats. Examples of such organizations include the Cyber Threat Alliance and the Information Sharing and Analysis Centers (ISACs), which span various industries. These organizations are crucial in enhancing collective cybersecurity resilience and promoting a collaborative approach to tackling cyber threats.

Threat Feed Type	Description
Deep and dark web	Threat research is a counterintelligence gathering effort in which security companies and researchers attempt to discover modern cyber adversaries' tactics, techniques, and procedures (TTPs). There are many companies and academic institutions engaged in primary cybersecurity research. Security solution providers with firewall and antimalware platforms derive much data from their customers' networks. As they assist customers with cybersecurity operations, they can analyze and publicize TTPs and their indicators. These organizations also operate honeynets to observe how hackers interact with vulnerable systems. The deep web and dark web are also sources of threat intelligence. The deep web is any part of the World Wide Web not indexed by a search engine. This includes pages that require registration, pages that block search indexing, unlinked pages, pages using nonstandard DNS, and content encoded nonstandardly. Within the deep web are areas that are deliberately concealed from "regular" browser access. • Dark net — is a network established as an overlay to internet infrastructure by software, such as The Onion Router (TOR), Freenet, or I2P, that acts to anonymize usage and prevent a third party from knowing about the existence of the network or analyzing any activity taking place over the network. Onion routing, for instance, uses multiple layers of encryption and relays between nodes to achieve this anonymity. • Dark web — websites, content, and services accessible only over a dark net. While there are dark web search engines, many sites are hidden from them. Access to a dark website via its URL is often only available via "word of mouth" bulletin boards. Investigating these dark websites and message boards is a valuable source of counterintelligence. The anonymity of dark web services has made it easy for investigators to infiltrate the forums and web stores set up to exchange stolen data and hacking tools. As adversaries react to this, they set up new networks and ways of identifying law enforcement infiltration. Consequently, dark nets and the dark web represent a continually shifting landscape. Please note that participating in illegal dark web activities is strictly prohibited. To stay safe, it is essential to exercise caution and follow legal and ethical guidelines when exploring the dark web. The dark web is generally associated with illicit activities and illegal content but has legitimate purposes. • Privacy and Anonymity — The dark web provides a platform for enhanced privacy and anonymity. It allows users to communicate and browse the internet without revealing their identity or location. It can be valuable for whistleblowers, journalists, activists, or individuals living under repressive government regimes. • Access to Censored Information — In countries with strict internet censorship, the dark web can be an avenue for accessing information that is otherwise blocked or restricted. It enables individuals to bypass censorship and access politically sensitive or controversial content. • o Research and Information Sharing — Some academic researchers or cybersecurity professionals may explore the dark web to gain insights into criminal activities and analyze emerging threats to improve cybersecurity operations.

Other Vulnerability Identification Methods

The following table reviews three additional vulnerability identification methods:

Method	Description
Penetration testing	Penetration testing, or pen testing, is a more aggressive approach to vulnerability management. Ethical hackers attempt to breach an organization's security in this practice, exploiting vulnerabilities to demonstrate their potential impact. While automated vulnerability scans and threat feeds are essential components of a robust security program, they may sometimes fail to identify specific vulnerabilities that a penetration test can uncover.
Bug bounties	Bug bounty programs are another proactive strategy and describe when organizations incentivize discovering and reporting vulnerabilities by rewarding external security researchers or "white hat" hackers. Both penetration testing and bug bounty programs are proactive cybersecurity practices to identify and mitigate vulnerabilities in a system or application. They both involve exploiting vulnerabilities to understand their potential impact, with the difference lying primarily in who conducts the testing and how it's structured. Penetration testing is typically performed by a hired team of professional ethical hackers within a confined time frame, using a structured approach based on the organization's requirements. This approach allows for a focused, in-depth examination of specific systems or applications and provides a predictable cost and timeline. In contrast, bug bounty programs open the testing process to a global community of independent security researchers. Rewards for finding and reporting vulnerabilities incentivize these researchers. This approach can bring diverse skills and perspectives to testing, potentially uncovering more complex or obscure vulnerabilities. An organization may choose penetration testing for a more controlled, targeted assessment, especially when testing specific components or meeting certain compliance requirements. A bug bounty program might be preferred when seeking a more extensive range of testing, leveraging the collective skills of a global community. However, many organizations see the value in both and use a combination of pen testing and bug bounty programs to ensure comprehensive vulnerability management. Responsible disclosure programs are established by organizations to encourage individuals to report security vulnerabilities in software or systems, allowing the organization to address and fix these vulnerabilities before they can be exploited maliciously. Responsible disclosure programs provide guidelines and procedures for reporting vulnerabilities and often reward or recognize individuals who responsibly disclose verified security issues.
Auditing	Auditing is an essential part of vulnerability management. Where product audits focus on specific features, such as application code, system/ process audits interrogate the wider use and deployment of products, including supply chain, configuration, support, monitoring, and cybersecurity. Security audits assess an organization's security controls, policies, and procedures, often using standards like ISO 27001 or the NIST Cybersecurity Framework as benchmarks. These audits can identify technical vulnerabilities and operational weaknesses impacting an organization's security posture. Cybersecurity audits are comprehensive reviews designed to ensure an organization's security posture aligns with established standards and best practices. There are various types of cybersecurity audits, including compliance audits, which assess adherence to regulations like GDPR or HIPAA; risk-based audits, which identify potential threats and vulnerabilities in an organization's systems and processes; and technical audits, which delve into the specifics of the organization's IT infrastructure, examining areas like network security, access controls, and data protection measures. Penetration testing fits into cybersecurity audit practices as a critical component of a technical audit as it provides a practical assessment of the organization's defenses by simulating real-world attack scenarios. Rather than simply evaluating policies or configurations, penetration tests seek exploitable vulnerabilities, providing a clear picture of what an attacker might achieve. The findings from these tests are then used to improve the organization's security controls and mitigate identified risks. Penetration tests also play an important role in compliance audits, as many regulations require organizations to conduct regular penetration testing as part of their cybersecurity program. For instance, the Payment Card Industry Data Security Standard (PCI DSS) mandates annual and proactive penetration tests for organizations handling cardholder data.

7.1.4 Vulnerability Analysis and Remediation (Lesson Video)

Transcript:

Vulnerability analysis involves evaluating vulnerabilities for their potential impact and exploitability. Vulnerability analysis aids in vulnerability classification, categorizing vulnerabilities based on their characteristics, such as the type of system or application affected, the nature of the exposure, or the potential impact.

Vulnerability analysis also helps to prioritize remediation. Remediation describes how vulnerabilities are addressed to mitigate their potential risk. Prioritization is typically based on factors such as vulnerability severity, the ease of exploitation, and the potential impact of an attack. Prioritizing vulnerabilities helps an organization focus limited resources on addressing the most significant threats first.

Mitigation techniques include applying patches, changing configurations, updating software, or replacing vulnerable systems. When immediate remediation is impossible, compensating controls describe alternative plans to reduce the risk.

Verification that remediation efforts have been successful is accomplished via several methods, including re-scanning affected systems. Organizations can significantly improve their resilience against cyberattacks by carefully analyzing and remediating vulnerabilities.

Vulnerability response and remediation practices encompass various strategies and tactics, including patching, insurance, segmentation, compensating controls, exceptions, and exemptions, each playing a distinct role in managing and mitigating cybersecurity risks.

Patching is one of the most straightforward and effective remediation practices. It involves applying updates and patches to software or systems to fix known vulnerabilities. Patching helps prevent attackers from exploiting known vulnerabilities, improving an organization's security posture.

Cybersecurity insurance can provide financial protection if a security breach results from a vulnerability. It's another factor in vulnerability response. While insurance doesn't mitigate vulnerabilities directly, it's vital in a comprehensive risk management strategy, complementing technical controls with financial risk transfer.

Segmentation involves dividing a network into separate segments to contain potential security breaches. If an attacker exploits a vulnerability and gains access to one segment, they're confined to it. This prevents them from moving laterally across the entire network, limiting the impact of a successful attack and supporting incident response teams.

Compensating controls refer to measures put in place to mitigate the risk of a vulnerability when security teams can't directly eliminate it or when direct remediation isn't immediately possible.

Exceptions and exemptions describe scenarios where specific vulnerabilities can't be remediated due to business criticality, technical constraints, or cost constraints. In these cases, the senior leadership teams accept the risk and document the rationale for the decision, along with an established timeline for reassessment.

Validating vulnerability remediation is critical for several vital reasons. Validation ensures that the remediation actions have been implemented correctly and function as intended. Despite best intentions, human error or technical problems can frequently lead to incomplete or incorrect implementation of fixes. These issues go unnoticed without validation, exposing the organization to the same vulnerability it initially sought to address.

Validation can be achieved through re-scanning, auditing, or other verification forms. Re-scanning involves performing additional vulnerability scans after remediation actions have been implemented. The re-scan aims to determine if the vulnerabilities identified in the initial scan have been resolved. If the same vulnerabilities aren't identified in the re-scan, it strongly indicates that the remediation efforts were successful.

Auditing involves an in-depth examination of the remediation process by reviewing the steps taken to address the vulnerability and ensuring they align with the organization's policies and best practices. Audits verify that necessary documentation has been updated, such as records of identified vulnerabilities, remediation actions taken, and any exceptions or exemptions granted.

Verification is the process of confirming the results of the remediation actions. It involves various methods, including manual checks, automated testing, or reviews of system logs or other records. Verification ensures that remediation steps have been implemented correctly, function as intended, and don't introduce new issues or vulnerabilities.

That's it for this lesson. In this lesson, we discussed vulnerability analysis and remediation. We also talked about the importance of validating vulnerability remediation to ensure the fixes were effective.

7.1.5 Vulnerability Analysis and Remediation Facts

This lesson covers the following topics:

- Vulnerability analysis
- Validating vulnerability remediation
- Vulnerability reporting

Vulnerability Analysis

Vulnerability analysis supports several critical aspects of an organization's cybersecurity strategy, including prioritization, vulnerability classification, exposure considerations, organizational impact, and risk tolerance contexts.

Note the following considerations related to vulnerability analysis:

Consideration	Description
Prioritization	Vulnerability analysis prioritizes remediation efforts by identifying the most critical vulnerabilities in an organization. Prioritization is typically based on factors such as vulnerability severity, the ease of exploitation, and the potential impact of an attack. Prioritizing vulnerabilities helps an organization focus limited resources on addressing the most significant threats first.
Classification	Vulnerability analysis aids in vulnerability classification, categorizing vulnerabilities based on their characteristics, such as the type of system or application affected, the nature of the exposure, or the potential impact. Classification can help clarify the scope and nature of an organization's threats.
Exposure factor	Vulnerability analysis must also consider exposure factors like the accessibility of a vulnerable system or data and environmental factors like the current threat landscape or the specifics of the organization's IT infrastructure. These factors can significantly influence the likelihood of a vulnerability being exploited and directly impact its overall risk level. The exposure factor (EF) represents the extent to which an asset is susceptible to being compromised or impacted by a specific vulnerability. It helps assess the potential impact or loss if the vulnerability is exploited. Factors might include weak authentication mechanisms, inadequate network segmentation, or insufficient access control methods.
Impacts	Vulnerability analysis assesses the potential organizational impact of vulnerabilities. This could include financial loss, reputational damage, operational disruption, or regulatory penalties. Understanding this impact is crucial for making informed decisions about risk mitigation.
Environmental factors	Several environmental variables play a significant role in influencing vulnerability analysis. One of the primary environmental factors is the organization's IT infrastructure, which includes the hardware, software, networks, and systems in use. These components' diversity, complexity, and age can affect the number and types of vulnerabilities present. For instance, legacy systems may have known unpatched vulnerabilities, while new emerging technologies might introduce unknown vulnerabilities. The external threat landscape is another crucial environmental factor. The prevalence of certain types of attacks or the activities of specific threat actors can affect the likelihood of exploitation of particular vulnerabilities. For example, if ransomware attacks rise within the medical industry, that sector can prioritize those vulnerabilities. The regulatory and compliance environment is another significant factor. Organizations in heavily

Consideration	Description
	regulated industries, like healthcare or finance, may need to prioritize vulnerabilities that could lead to sensitive data breaches and result in regulatory penalties. The operational environment, including the organization's workflows, business processes, and usage patterns, can also influence vulnerability analysis. Certain operational practices increase exposure to specific vulnerabilities or affect the potential impact of a successful exploit. Examples include poor patch management practices, lack of rigorous access controls, lack of awareness training, poor configuration management practices, and insufficient application development policies.
Risk tolerance	Vulnerability analysis must align with an organization's risk tolerance. Risk tolerance refers to the level of risk an organization is willing to accept, which can vary greatly depending on the organization's size, industry, regulatory environment, and strategic objectives. By aligning vulnerability analysis with risk tolerance, an organization can ensure its vulnerability management efforts align with its overall risk management strategy.

Validating Vulnerability Remediation

Validating vulnerability remediation is critically essential for several key reasons. Validation ensures that the remediation actions have been implemented correctly and function as intended. Despite best intentions, human error or technical problems can frequently lead to incomplete or incorrect implementation of fixes. These issues go unnoticed without validation, exposing the organization to the same vulnerability it initially sought to address.

Validation helps confirm that the remediation has not inadvertently introduced new issues or vulnerabilities. For example, a patch may interfere with other software or systems, or a configuration change could expose new security gaps. Also, validation provides a measure of accountability, ensuring that responsible parties adequately address identified vulnerabilities. This is especially important in larger organizations where multiple teams or individuals may be involved in the remediation process.

- Re-scanning involves performing additional vulnerability scans after remediation actions have been implemented. The re-scan aims to determine if the vulnerabilities identified in the initial scan have been resolved. If the same vulnerabilities are not identified in the re-scan, it strongly indicates that the remediation efforts were successful.
- Auditing involves an in-depth examination of the remediation process by reviewing the steps taken to address the vulnerability and ensuring they align with the organization's policies and best practices. Audits verify that necessary documentation has been updated, such as records of identified vulnerabilities, remediation actions taken, and any exceptions or exemptions granted.
- Verification confirms the results of the remediation actions. It involves various methods, including manual checks, automated testing, or reviews of system logs or other records. Verification ensures that remediation steps have been implemented correctly, function as intended, and do not introduce new issues or vulnerabilities.

Vulnerability Reporting

Vulnerability reporting is a crucial aspect of vulnerability management and is critical in maintaining an organization's cybersecurity posture. A comprehensive vulnerability report highlights the existing vulnerabilities. It ranks them based on their severity and potential impact on the organization's assets, enabling the management to prioritize remediation efforts effectively.

The Common Vulnerability Scoring System (CVSS) provides a standardized method for rating the severity of vulnerabilities. It includes metrics such as exploitability, impact, and remediation level. By using CVSS, organizations can compare and prioritize vulnerabilities consistently.

Another essential practice is to include information about the potential impact of each vulnerability in the report. This could involve describing the possible outcomes of exploiting the vulnerability, including data breaches, system outages, or other operational impacts. It is essential to provide recommendations for addressing each vulnerability in the report. Recommendations might suggest specific patches or updates, recommend configuration changes, or identify other mitigation strategies.

Timely reporting is also essential, as delays in reporting can lead to delays in remediation and increase the window of opportunity for attackers. Vulnerability reports must use a clear, concise format that is easy for technical and non-technical stakeholders to understand to help ensure that the report is understood and that appropriate actions are taken in response.

7.1.6 Explore End of Life Software / Hardware (Demo Video)

Transcript:

End-of-life in regard to software means there's a date set when it will no longer receive any updates. This normally means there's a new version already out or going to be released, where developers will focus on the new version of the software versus the old. We're going to explore some ways to keep track of the end-of-life.

One good way to keep track of this is by finding information on the internet. When a piece of software will no longer be supported, the date will be published so tech administrators will know when an update is needed. One website we can check out is endoflife.date. As you can see, all sorts of different software companies are listed. Some of the most popular pages are listed here in the middle. Let's take a look at Windows Server.
After scrolling down, we can see that Windows Server 2012 R2 will reach the end-of-life starting on October 10, 2023. Although this version of Windows Server hasn't received full support since 2018, it has been getting regular security updates. If a software isn't receiving security updates, it could leave you vulnerable to hackers. That's why, in most cases, software will be listed as end-of-life when it no longer receives any updates. Now, the last window is what they call extended security updates. Usually, this involves paying a substantial amount of money to receive updates past the normal security support date. Once the extended date is reached, there are typically no more updates to that version of the software.
Let's look at one more example. Red Hat Enterprise Linux goes through a similar phase. There comes a time when full support isn't available. Full Support may entail bug fixes and features, while Maintenance Support typically includes security updates and some bug fixes. RHEL 7 will reach the end-of-life starting on June 30th, 2024. There's also an Extended Life Cycle Support up through June 30th, 2028. If you work in the tech industry, you should always be aware of the end-of-life dates set by the software you manage.
Another way to be informed about end-of-life software on your network is to scan your computers. Network scanners can scan computers and find what software may be installed on them. Here, you can see that the CentOS 6.9 server is a device that was scanned on our network; however, it's at the end-of-life since it's not supported by regular maintenance. Instead of going to each computer and looking at a list, this could be helpful. The same concept applies to software. Installed software could also be marked as end-of-life if it's no longer supported. This could mean there's a new version of the software out, or the developer who originally made it abandoned the project. One example is PuTTY. You can see it picked up version 0.79. If this same version were on other systems, you'd see multiple computers listed here. The version that's installed can be used as a comparison to the latest version released by the developer. Knowing what kind of operating systems and software are installed on your network is a significant step in keeping your network secure. That's it for this demo. In this demo, we showed you how to research end-of-life dates and how you can use a network scanner to discover information on your network.

7.1.7 Identify Types of Vulnerabilities

7.1.8 Practice Questions (Section Quiz)

q_vuln_type_android_sec8

You are a cybersecurity analyst at a tech company that develops mobile applications.

Your team has been informed of a significant vulnerability that affects mobile operating systems. The vulnerability allows unauthorized remote access to sensitive user data.

Which of the following operating systems should your team prioritize for patching and security updates, based on the information provided?

Answers:

- iOS
- Windows Mobile
- BlackBerry OS
- ***Android**

Explanation:

Android is the correct answer. Given the information provided in the question, the vulnerability affects mobile operating systems and Android is one of the most widely used mobile operating systems in the world. Therefore, prioritizing Android for patching and security updates would be the most effective strategy to protect the largest number of users.

iOS is a widely used mobile operating system, and while it's important to address vulnerabilities in iOS, the question does not specify that iOS is affected by this particular vulnerability. Therefore, prioritizing iOS would not be the most effective use of resources in this scenario.

Windows Mobile, while still in use in some places, has a much smaller market share compared to Android. Prioritizing Windows Mobile over Android, which has a larger user base, would not be the most effective strategy.

BlackBerry OS, like Windows Mobile, has a smaller market share and is not as widely used as Android. Therefore, prioritizing BlackBerry OS would not be the most effective use of resources.

q_vuln_type_eol_sec8

You are the IT Director at a mid-sized company. The company's core business operations rely on a legacy system that has recently reached its end-of-life (EOL).

The vendor no longer supports this system, and it has known security vulnerabilities. However, transitioning to a new system would be costly and time-consuming.

What is the BEST course of action?

Answers:

- Continue using the EOL system and hope that the known vulnerabilities won't be exploited.
- Immediately transition to a new system, regardless of the cost and disruption to business operations.
- ***Develop a detailed transition plan that includes a cost-benefit analysis, risk assessment, and timeline for transitioning to a new system.**
- Ignore the EOL status and continue operations as usual, as the system has been reliable so far.

Explanation:

Developing a detailed transition plan that includes a cost-benefit analysis, risk assessment, and timeline for transitioning to a new system is the best option. Developing a detailed transition plan allows the company to understand the costs, benefits, and risks associated with transitioning to a new system. This approach ensures that the company can make an informed decision and plan for a smooth transition that minimizes disruption to business operations.

Continuing to use the EOL system and hope that the known vulnerabilities won't be exploited is not advisable. Continuing to use an EOL system with known vulnerabilities exposes the company to significant security risks.

Immediately transitioning to a new system, regardless of the cost and disruption to business operations, while addressing the security concern, may not be the best course of action. An immediate transition without proper planning could lead to unforeseen costs, potential data loss, and disruption to business operations.

Ignoring the EOL status and continue operations as usual, as the system has been reliable so far, is not advisable. Ignoring the EOL status of a system is a risky approach. Even if the system has been reliable so far, the lack of vendor support and potential security vulnerabilities pose significant risks to the company.

q_vuln_type_firmware_sec8

A cyber consultant inspects a corporate desktop after receiving numerous complaints.

What type of vulnerability can include instances where processors inside the computer allow malicious programs to steal data during processing?

Answers:

- ***Firmware**
- Virtualization
- End-of-life
- Legacy

Explanation:

Firmware vulnerabilities include instances where processors inside the computer allow malicious programs to steal data during processing.

Virtualization vulnerabilities include virtual machine (VM) escape--when an attacker with access to a VM breaks out of this isolated environment and gains access to the host system or other VMs running on the same host.

An end-of-life (EOL) system vulnerability includes instances where a specific product or version of a product that the manufacturer or vendor publicly declares as no longer supported.

Legacy system vulnerabilities typically describe outdated software methods, technology, computer systems, or application programs with continued use despite known shortcomings.

q_vuln_type_legacy_01_sec8

A software engineer at a growing tech company identifies that some divisions in the organization still operate on legacy systems. The firmware for these systems has not seen updates in over a decade.

The chief information security officer (CISO) recognizes the imminent risks these outdated systems pose and decides to hold a training session. During the training, the software engineer asks about the main vulnerability of such systems.

Given the context of legacy and end-of-life system vulnerabilities, what is the primary risk of using firmware that has not received security updates, thus potentially exposing the system to breaches?

Answers:

- The firmware becomes faster and more efficient.
- The system can experience compatibility issues with newer software.
- ***Unauthorized access becomes easier for potential attackers.**
- The system may require frequent restarts due to firmware instability.

Explanation:

Non-updated legacy firmware remains unprotected from patches designed to rectify known vulnerabilities, making it an inviting target for attackers who can exploit these weak spots to achieve unauthorized system access.

Outdated firmware does not magically enhance its performance over time. In fact, as technology advances, it tends to become less compatible and loses efficiency when juxtaposed against its updated counterparts.

Compatibility issues pose challenges for legacy systems; however, when evaluated from a security standpoint, the foremost risk they present is the heightened susceptibility to potential security breaches.

While occasional system instability might contribute to aging firmware, this inconsistency does not measure up to the predominant security threat that outdated firmware poses regarding potential unauthorized penetrations.

q_vuln_type_legacy_02_sec8

What describes systems that create risk because they no longer receive critical security updates and patches?

Answers:

- ***Legacy system**
- Operating system
- Sandbox
- Thin client

Explanation:

A legacy system is an outdated computing software or hardware still in use. Legacy systems generally receive no support or maintenance.

An operating system is software that manages hardware while providing services for applications.

A sandbox is an isolated area created for testing and developing host and system environments. Administrators can use a sandbox to test patches and security controls before production.

A thin client is a low-power computer that runs from resources stored on a central server. A thin client connects remotely to a server-based computing environment that stores applications, data, and memory.

q_vuln_type_legacy_03_sec8

A cyber technician works on a corporate laptop where an employee complains the software is outdated.

What type of vulnerability describes continued use of outdated software methods, technology, computer systems, or application programs, despite known shortcomings?

Answers:

- ***Legacy**
- Firmware
- Virtualization
- End-of-life

Explanation:

Legacy system vulnerabilities typically describe the use of outdated software methods, technology, computer systems, or application programs despite known shortcomings.

Firmware vulnerabilities include instances where processors inside the computer allow malicious programs to steal data during processing.

Virtualization vulnerabilities include virtual machine (VM) escape--when an attacker with access to a virtual machine breaks out of this isolated environment and gains access to the host system or other VMs running on the same host.

An end-of-life (EOL) system vulnerability includes instances where a specific product or version of a product that the manufacturer or vendor has publicly declared as no longer supported.

q_vuln_type_os_01_sec8

You are the IT security manager at a large corporation. Your team has just discovered a significant vulnerability in the company's Linux-based server infrastructure. The vulnerability, if exploited, could allow an attacker to gain unauthorized access to sensitive data.

Your team has identified a patch that can fix the vulnerability, but applying the patch will require significant downtime during peak business hours.

What is the BEST course of action?

Answers:

- Apply the patch immediately during peak business hours to ensure the vulnerability is fixed as soon as possible.
- Wait until off-peak hours to apply the patch to minimize business disruption.
- Ignore the patch since Linux is known for its security and the likelihood of an attack is low.
- ***Inform the senior management about the vulnerability and the potential impact of the patch on business operations, and ask for their decision.**

Explanation:

Informing the senior management about the vulnerability and the potential impact of the patch on business operations is the best option. As the IT security manager, it's your responsibility to inform senior management about the vulnerability, the potential impact of an attack, and the potential impact of the patch on business operations. This allows the senior management to make an informed decision about how to proceed, taking into consideration both the security risk and the potential business impact.

Applying the patch immediately during peak business hours to ensure the vulnerability is fixed as soon as possible is not the best because it could disrupt business operations during peak hours. While it's important to fix vulnerabilities as soon as possible, it's also important to consider the impact on the business.

Waiting until off-peak hours to apply the patch to minimize business disruption might seem reasonable, but waiting to apply the patch could leave the system vulnerable to attacks. If an attacker were to exploit the vulnerability before the patch is applied, the consequences could be severe.

Ignoring the patch since Linux is known for its security and the likelihood of an attack is low is not advisable. Ignoring a known vulnerability, especially one that could allow unauthorized access to sensitive data, is not a good security practice, regardless of the perceived security of the operating system.

q_vuln_type_os_02_sec8

A large multinational corporation has recently discovered a significant vulnerability in their widely used operating system. The vulnerability could potentially allow unauthorized remote access to sensitive corporate data.

The corporation's IT team has been tasked with addressing this issue.

Which of the following approaches would be the most effective in managing this vulnerability?

Answers:

- Ignore the vulnerability since the operating system is due to be updated in the next six months.
- Inform all employees about the vulnerability and ask them to be extra vigilant.
- ***Implement a patch to fix the vulnerability and conduct a thorough system-wide security audit.**
- Disconnect all systems from the network until a new operating system can be installed.

Explanation:

Implementing a patch to fix the vulnerability is the most effective immediate response. This will close the security gap and protect the corporation's data. Conducting a thorough system-wide security audit will help identify any other potential vulnerabilities and ensure that the patch has been implemented correctly across all systems.

Ignoring the vulnerability is not a viable option. Even though the operating system is due to be updated, six months is a long time for a significant vulnerability to be left unaddressed. During this time, the corporation is at risk of a potential data breach, which could have severe financial and reputational consequences.

While it's important to inform employees about potential security threats, simply asking them to be extra vigilant is not enough. Employees may not have the technical knowledge or ability to protect their systems from a significant operating system vulnerability.

Disconnecting all systems from the network is an extreme measure that would likely disrupt business operations. While this might prevent external unauthorized access, it's not a practical or effective long-term solution. It also doesn't address the vulnerability itself, which would still be present when the systems are reconnected to the network.

q_vuln_type_prioritize_sec8

The security team at a major corporation has discovered multiple vulnerabilities during its latest assessment. The security manager must prioritize these vulnerabilities to ensure that the most critical ones get addressed first.

In the context of vulnerability management and prioritization, which of the following criteria are MOST crucial for the security manager to consider when determining the urgency of addressing a specific vulnerability?

Answers:

- ***The potential impact of the vulnerability on the organization's core operations**

- The age of the software containing the vulnerability
- The number of times the vulnerability has appeared in past assessments
- The popularity of the software among the organization's employees

Explanation:

Organizations primarily consider the potential operational impact when prioritizing vulnerabilities. They should urgently address exploitations that could cause severe operational disruptions, data breaches, or other significant problems.

Older software may be more vulnerable, but its age does not primarily dictate prioritization. Instead, the vulnerability's severity and potential impact hold greater importance.

Recurring vulnerabilities might highlight systemic problems, but the frequency of past identifications does not determine the vulnerability's criticality.

Although the popularity of software can show the extent of a vulnerability within an organization, it does not directly signify its severity or potential operational impact.

q_vuln_assess_bug_bounty_secp8

A medium-sized software development company recently introduced a bug bounty program to identify and mitigate vulnerabilities in their flagship application. The security manager plans to coordinate the program's rules and engagement policies.

When setting up a bug bounty program for vulnerability management, which activities should the security manager prioritize to ensure the program's effectiveness and ethical participation? (Select two.)

Answers:

- ***Establishing a clear scope of which assets researchers can test.**
- Offering substantial rewards regardless of the severity of the bug found.
- ***Providing a secure platform for researchers to report findings.**
- Allowing researchers to disclose findings publicly immediately after discovery.
- Providing valuable real-time information on the latest cyber threats and vulnerabilities.

Explanation:

Security managers must define the scope to guide researchers about which areas they can and cannot test, preventing unintended system breaches outside the bug bounty program's intention.

A bug bounty program requires a secure and straightforward method for researchers to report vulnerabilities effectively.

The program should set rewards based on the bug's severity and impact. Offering high rewards for minor issues might result in financial inefficiencies and not attract top-notch researchers.

Disclosing vulnerabilities publicly before applying a patch compromises the software's security. Well-structured bug bounty programs mandate researchers to delay public disclosures until addressing the vulnerability.

Open-source and proprietary threat feeds provide valuable real-time information on the latest cyber threats and vulnerabilities.

q_vuln_assess_dark_web_secp8

As a cybersecurity analyst, you are investigating a case of corporate espionage. You suspect that the perpetrator may have used the deep or dark web to access confidential information.

Which of the following would be a legitimate reason for someone to use the deep or dark web in a corporate setting?

Answers:

- To conduct illegal activities anonymously
- To bypass corporate network restrictions
- ***To access censored information**
- To avoid detection by law enforcement

Explanation:

To access censored information is the correct answer. In some situations, the deep or dark web can be used legitimately to access information that has been censored or is otherwise difficult to obtain. For example, a cybersecurity analyst might use the deep web to access information about emerging threats that is not available through traditional channels.

While the deep and dark web can provide anonymity, using it for illegal activities is not a legitimate reason in a corporate setting. Engaging in illegal activities can lead to severe consequences, including legal penalties and damage to the company's reputation.

While it's technically possible to use the deep and dark web to bypass network restrictions, this is not a legitimate or ethical use in a corporate setting. Network restrictions are typically in place for a reason, such as to protect the company's network and data.

While the deep and dark web can provide a degree of anonymity, using it to avoid detection by law enforcement is not a legitimate reason in a corporate setting. This could imply illegal activities, which are not acceptable in a corporate environment.

q_vuln_assess_disclosure_program_secp8

A system admin discovers a security vulnerability in a widely used software and brings it to the manager's attention, who said to fix it but has not released the information yet. The system admin releases the information anyway, as the company is part of a voluntary info-sharing organization.

What is the company a part of?

Answers:

- ***Responsible disclosure program**
- Bug bounty program
- Penetration testing
- Auditing

Explanation:

An organization's responsible disclosure program encourage individuals to report security vulnerabilities in software or systems, allowing organizations to address and fix these vulnerabilities before exploitation.

Corporations that sell software utilize a bug bounty program. They offer cash incentives to reveal vulnerabilities to them before the public so that the corporation can patch the vulnerability.

Penetration testing against a company's network ensures vulnerabilities are not readily exploitable. It does not deal with releasing information relating to vulnerabilities.

Auditing assesses an organization's security controls, policies, and procedures, often using standards like ISO 27001 or the NIST Cybersecurity Framework as benchmarks. They ensure a network's security against vulnerabilities unrelated to releasing information on vulnerabilities.

q_vuln_assess_osint_secp8

An IT administrator has reviewed security policies and best practices on well-known IT bulletin boards. During this reading, the admin became aware of several new vulnerability exploits.

What type of information is this?

Answers:

- ***Open-source intelligence**
- Information-sharing organizations
- Cyber threat intelligence
- Threat feeds

Explanation:

Open-source intelligence (OSINT) describes collecting and analyzing publicly available information and using it to support decision-making. OSINT identifies vulnerabilities and threat information by gathering data from many sources such as blogs, forums, social media platforms, and even the dark web.

Information-sharing organizations are groups composed of businesses, government entities, and academic institutions. Each member of an organization benefits from sharing information with the group members.

Cyber threat intelligence is investigating, collecting, analyzing, and disseminating information about emerging threats and threat sources.

Threat feeds are signatures and pattern-matching rules supplied to analysis platforms as an automated feed. Companies respond swifter to emergent threats, as threat feeds provide real-time information.

q_vuln_assess_penetration_testing_secp8

A large financial institution is considering outsourcing its IT infrastructure to a third-party cloud service provider. The company has concerns about the risks of giving its sensitive financial data to an external vendor.

What approach should the company use to ensure the vendor complies with the appropriate security standards and regulations?

Answers:

- Enter into a contract without clauses for regular assessments or audits of the vendor's security practices.
- Rely on the vendor's reputation in the industry without conducting any further assessments.
- ***Conduct penetration testing on the vendor's infrastructure or seek evidence that the vendor has performed regular penetration tests.**
- Prioritize the vendor's cost and ease of use over security considerations.

Explanation:

Penetration testing identifies potential vulnerabilities in a vendor's systems, networks, and applications, assessing their security posture. The company gains insights into the vendor's security resilience and vulnerabilities that attackers could exploit by conducting or requesting evidence of regular penetration tests.

Simply entering a contract without provisions for regular assessments or audits may be insufficient. Regular assessments and audits ensure ongoing adherence to security standards and compliance requirements.

Solely relying on the vendor's reputation may not offer a comprehensive and objective assessment of the vendor's current security practices and regulatory compliance.

While important, cost and ease of use should not take precedence over security considerations.

q_vuln_assess_proprietary_threat_feeds_secp8

Open-source threat feeds are an excellent tool for utilizing all companies with an online presence.

However, some companies use proprietary threat feeds for an additional cost due to more depth, breadth, and sophistication of analysis found herein.

What are the three primary forms these can take? (Select three.)

Answers:

- ***Behavioral threat research**
- ***Reputational threat intelligence**
- ***Threat data**
- Vulnerability management
- Dark net
- Dark web
- Bug bounties

Explanation:

The following are the three primary forms of proprietary threat feeds:

- Behavioral threat research is a narrative commentary describing examples of attacks, tactics, techniques, and procedures (TTPs) gathered through primary research sources.
- Reputational threat intelligence is lists of Internet Protocol (IP) addresses and domains associated with malicious behavior, plus signatures of known file-based malware.
- Threat data is computer data that correlates events observed on a customer's networks and logs with known TTPs, and threat actor indicators.

The following are NOT included as primary forms of proprietary threat feeds:

- Vulnerability management is a cornerstone of modern cybersecurity practices aimed at identifying, classifying, remediating, and mitigating vulnerabilities within a system or network.
- The dark net and the dark web are deep web areas that are deliberately concealed from "regular" browser access.
- Bug bounties are related to bug bounty programs that open the testing process to a global community of independent security researchers.

q_vuln_assess_third-party_threat_feed_secp8

As the new chief information security officer (CISO) for a mid-sized company, you are tasked with enhancing the organization's vulnerability management practices.

You have a limited budget but need to ensure that your organization stays ahead of the latest threats.

Which type of threat feed would be the MOST effective choice for your situation?

Answers:

- Open-source intelligence
- Information-sharing organization
- ***Third-party threat feeds**
- Internal threat feeds

Explanation:

Third-party threat feeds is the correct answer. Third-party threat feeds, both open-source and proprietary, provide real-time information on the latest cyber threats and vulnerabilities. They aggregate data from various sources and can be integrated into an organization's security infrastructure. This makes them a cost-effective solution for a mid-sized company with a limited budget, as they provide comprehensive coverage and advanced analytic insights.

While open-source intelligence (OSINT) can provide valuable information about potential threats and vulnerabilities, it requires significant time and resources to collect, analyze, and apply this information effectively. It might not be the most efficient choice for a mid-sized company with a limited budget.

Information-sharing organizations can provide valuable threat intelligence, but their effectiveness can vary depending on the members' contributions and the organization's resources. They might not provide the real-time, comprehensive coverage that a mid-sized company needs to stay ahead of threats.

While internal threat feeds can provide valuable insights based on the organization's specific experiences, they might not provide the broad coverage necessary to stay ahead of the latest threats. They also require significant resources to maintain and analyze, which might not be feasible for a mid-sized company with a limited budget.

q_vuln_assess_threat_secp8

In your role as a security analyst, you need to stay up to date on the latest threats. You are currently reviewing the latest real-time updates on cyberthreats from across the world.

Which of the following resources are you MOST likely using?

Answers:

- ***Threat feeds**
- Advisories and bulletins
- Intelligence fusion
- Threat hunting

Explanation:

Threat feeds provide real-time updates on cyberthreats across the world. They can provide information such as suspicious domains, known malware, known malicious IP addresses, and more.

Advisories and bulletins are usually updated weekly and provide much more detailed information on the newest threats.

Intelligence fusion is the sharing of information between multiple government agencies and private security firms.

Threat hunting is the human-based, proactive and methodical monitoring of a network, systems, and software. This is done in order to detect any suspicious activity that may have evaded the automated tools.

q_vuln_assess_vuln_01_secp8

You want to be able to identify the services running on a set of servers on your network.

Which tool would BEST give you the information you need?

Answers:

- Port scanner
- Protocol analyzer
- ***Vulnerability scanner**
- Network mapper

Explanation:

Use a vulnerability scanner to gather information about systems such as the applications or services running on a system. A vulnerability scanner often combines functions found in other tools and can perform additional functions, such as identifying open firewall ports, missing patches, and default or blank passwords.

A port scanner is a tool that probes systems for open ports. A port scanner tells you which ports are opened in the firewall, but it cannot identify services running on a server if the firewall port has been closed. A network mapper is a tool that can discover devices on a network and shows those devices in a graphical representation. Network mappers typically use a ping scan to discover devices and a port scanner to identify open ports on those devices.

Use a protocol analyzer to identify traffic that is sent on the network medium and traffic sources. Services could still be running on a server that do not generate network traffic that a protocol analyzer would catch.

q_vuln_assess_vuln_02_secp8

You have run a vulnerability scanning tool and identified several patches that need to be applied to a system.

What should you do next after applying the patches?

Answers:

- ***Run the vulnerability assessment again.**
- Use a port scanner to check for open ports.
- Update the vulnerability scanner definition files.
- Document your actions.

Explanation:

After fixing an identified vulnerability, you should re-run the vulnerability scan to verify that everything has been fixed and that additional issues are not present.

You should update definition files before you run the first scan. Using a port scanner is unnecessary because most vulnerability scanners include a check of open ports. Documenting your actions should occur after you have finished all necessary actions.

q_vuln_analysis_analysis_secp8

A cyber engineer enhances processes and controls surrounding exposures and vulnerabilities to meet all regulatory requirements before a year-end inspection.

Which of the following focuses on key aspects of the organization's cybersecurity strategy, including prioritization, considerations of exposure, and risk tolerance contexts?

Answers:

- ***Vulnerability analysis**
- Open-source intelligence (OSINT)
- Common Vulnerability Scoring System (CVSS)
- Common Vulnerabilities and Exposures (CVE)

Explanation:

Vulnerability analysis is critical in supporting several key aspects of an organization's cybersecurity strategy, including prioritization, vulnerability classification, exposure considerations, organizational impact, and risk tolerance contexts.

Not best suited for this scenario, open-source intelligence describes collecting and analyzing publicly available information and using it to support decision-making.

Additionally, the Common Vulnerability Scoring System (CVSS) is the product of the Forum of Incident Response and Security Teams, and its metrics generate a score from 0 to 10 based on vulnerability characteristics.

Common Vulnerabilities and Exposures (CVE) is a dictionary of vulnerabilities in published operating systems and applications software and provides the principal input for NIST's National Vulnerability Database.

q_vuln_analysis_auditing_secp8

You are a cybersecurity analyst who has recently implemented a remediation plan for a critical vulnerability in your organization's network. Your manager has asked you to ensure the effectiveness of your remediation efforts.

What is the MOST crucial next step you should take?

Answers:

- Inform the stakeholders about the remediation
- Update the organization's vulnerability database
- ***Auditing the remediation process**
- Move on to the next identified vulnerability

Explanation:

Auditing the remediation process is the correct answer. Auditing involves an in-depth examination of the remediation process by reviewing the steps taken to address the vulnerability and ensuring they align with the organization's policies and best practices. Audits verify that necessary documentation has been updated, such as records of identified vulnerabilities, remediation actions taken, and any exceptions or exemptions granted. This step is crucial in validating the effectiveness of the remediation efforts.

While it's important to keep stakeholders informed, it's not the most crucial next step. Before informing stakeholders, you need to ensure that the remediation has been successful and hasn't introduced new vulnerabilities.

While updating the vulnerability database is an important step in the vulnerability management process, it should be done after validating the remediation. If the remediation was not successful or introduced new vulnerabilities, the database would need to be updated accordingly.

Moving on to the next vulnerability without validating the remediation of the current one could leave your organization exposed. If the remediation was not successful or introduced new vulnerabilities, you would not be aware of this without validation.

q_vuln_analysis_exposure_secp8

You are the lead cybersecurity analyst for a multinational corporation. Your team has recently completed a vulnerability analysis of the company's IT infrastructure.

The CEO has requested a briefing on the most critical consideration that should guide the company's remediation efforts.

Which of the following considerations should you emphasize in your briefing?

Answers:

- Prioritization
- Classification
- ***Exposure factor**
- Risk tolerance

Explanation:

Exposure factor is the correct answer. The exposure factor represents the extent to which an asset is susceptible to being compromised or impacted by a specific vulnerability. It considers the accessibility of a vulnerable system or data and environmental factors like the current threat landscape or the specifics of the organization's IT infrastructure. These factors significantly influence the likelihood of a vulnerability being exploited and directly impact its overall risk level. Therefore, understanding the exposure factor is crucial in guiding remediation efforts.

While prioritization is an important aspect of vulnerability analysis, it is not the most critical consideration. Prioritization helps to identify the most critical vulnerabilities based on factors such as severity, ease of exploitation, and potential impact. However, without considering the exposure factor, prioritization might not accurately reflect the actual risk to the organization.

Classification aids in categorizing vulnerabilities based on their characteristics. While it helps to clarify the scope and nature of threats, it does not directly influence the remediation efforts. It is a step in the process, but not the most critical consideration.

Risk tolerance refers to the level of risk an organization is willing to accept. While it is important to align vulnerability analysis with risk tolerance, it is not the most critical consideration. Risk tolerance is more about the organization's overall risk management strategy, while the exposure factor directly influences the immediate remediation efforts.

q_vuln_analysis_prioritization_secp8

A company has added several new assets and software to its system and is meeting to review its risk matrix. It wants to ensure risk management efforts focus on vulnerabilities MOST likely impacting its operations significantly.

What is this commonly referred to as?

Answers:

- ***Prioritization**
- Risk tolerance
- Classification
- Environmental variables

Explanation:

Prioritization of vulnerabilities that affect a company ensures the resolution of the most critical vulnerabilities first. The measure of the impact and exposure of the vulnerabilities determines the rank on prioritization.

Risk tolerance is the level of risk that a company is willing to accept in its daily operations, not the ranking of risks.

Classification categorizes vulnerabilities based on their characteristics, such as the type of system or application affected, the nature of the vulnerability, or the potential impact.

Environmental variables related to vulnerability analysis refer to the systems and software a company has currently in use. Age and diversity, as well as complexity, are all variables.

q_vuln_analysis_reporting_secp8

You are a cybersecurity analyst for a large corporation. Your team has recently completed a vulnerability analysis, and you are tasked with creating a comprehensive vulnerability report.

What is the MOST critical information to include in your report to effectively guide the organization's remediation efforts?

Answers:

- The names of the team members who conducted the vulnerability analysis.
- The software tools used in the vulnerability analysis.
- ***The potential impact of each vulnerability and recommendations for addressing them.**
- The number of vulnerabilities found in each department.

Explanation:

The potential impact of each vulnerability and recommendations for addressing them is the correct answer. This information is crucial for helping the organization understand the severity of each vulnerability and how to address it. Including the potential impact helps the organization prioritize which vulnerabilities to address first, and providing recommendations for addressing each vulnerability gives the organization a clear path forward.

While it's important to acknowledge the work of the team, this information does not directly contribute to the organization's remediation efforts and is therefore not the most critical information to include in the report.

While it's good to document the tools used for transparency and reproducibility, this information is not the most critical for guiding the organization's remediation efforts.

While this information can provide some insights into where vulnerabilities are concentrated, it does not provide enough detail to guide the organization's remediation efforts effectively. It's more important to understand the potential impact of each vulnerability and how to address it.

7.2 Vulnerability Scanning

As you study this section, answer the following questions:

- What is the purpose of vulnerability scanning?
- What information is critical in vulnerability scanning?
- What software is available to complete a vulnerability scan on a network?

In this section, you will learn to:

- Conduct vulnerability scans.
- Scan a network with Nessus.
- Scan a network with OpenVAS.
- Scan for cleartext vulnerabilities.
- Scan for FTP vulnerabilities.
- Scan for TLS vulnerabilities.
- Scan for Windows vulnerabilities.
- Scan for Linux vulnerabilities.
- Scan for domain controller vulnerabilities.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.2 Explain common threat vectors and attack surfaces.
	• Vulnerable software
	
	4.3 Explain various activities associated with vulnerability management.
	• Identification methods
	  ○ Vulnerability scan
	  ○ Application security
	  ▪ Static analysis
	  ▪ Dynamic analysis
	  ▪ Package monitoring
	• Analysis
	  ○ Common Vulnerability Enumeration (CVE)
	  ○ Vulnerability classification
	
	4.4 Explain security alerting and monitoring concepts and tools.
	• Monitoring computing resources
	  ○ Systems
	  ○ Applications
	  ○ Infrastructure

	• Activities
	  ○ Scanning
	  ○ Reporting
	• Tools
	  ○ Vulnerability scanners
	
	4.9 Given a scenario, use data sources to support an investigation.
	
	• Data sources
	  ○ Dashboards

TestOut Security Pro	5.2 Assessment techniques 5.2.3 Scan for vulnerabilities

7.2.1 Vulnerability Scanning (Lesson Video)

Transcript:

A vulnerability scan is a security assessment technique used to identify and evaluate potential weaknesses or vulnerabilities in a computer system, network, or application. It involves scanning the target system for known vulnerabilities, misconfigurations, or security loopholes that could be exploited by attackers. Vulnerability scans are performed internally and externally to inventory vulnerabilities from different network viewpoints. Vulnerabilities identified during scanning are then classified and prioritized for remediation by security operations teams.

A network vulnerability scanner, such as Tenable Nessus or OpenVAS, is designed to test network hosts, including client PCs, mobile devices, servers, routers, and switches. It performs three primary functions: discovery, assessment, and prioritization. During discovery, it examines an organization's on-premises systems, applications, and devices. It then assesses the scan results, comparing them to configuration templates and lists of known vulnerabilities.
Typical results from a vulnerability assessment will identify missing patches, deviations from baseline configuration templates, and other related vulnerabilities. The tool compiles a report about each vulnerability. Each identified vulnerability is categorized and prioritized using an assigned impact warning. Most tools also suggest remediation techniques. This information is highly sensitive, so using these tools and distributing the reports produced should be restricted to authorized hosts and user accounts.
Credentialed and non-credentialed scans are two different approaches to conducting vulnerability scans.
A credentialed scan, also known as an authenticated scan, requires the scanning tool to provide valid credentials to access the target system. This allows the scan to gather more comprehensive information about the system's configuration, installed software, and other details that may not be accessible without proper authentication.
In contrast, non-credentialed scans, also known as unauthenticated scans, don't require credentials to access the target system. Instead, it relies on external scanning techniques to probe for vulnerabilities and potential security issues.
Both types of scans have their own advantages and limitations. Credentialed scans offer a more accurate assessment of vulnerabilities but require privileged access to the target system. On the other hand, non-credentialed scans can be performed remotely and quickly, making them suitable for high-level assessments or initial security checks. However, they generally provide a less detailed view of the system's security posture compared to the credentialed scans.
Similarly, application vulnerability scanning describes a specialized method for identifying software application weaknesses. This includes static analysis and dynamic analysis. Static analysis involves reviewing application code without executing it. It involves examining the code for potential vulnerabilities, bugs, coding standards violations, and other issues that could impact the security, performance, or maintainability of the software.
Dynamic analysis involves executing a program and observing its behavior in real-time. It also tests running applications that can identify issues like unvalidated inputs, broken access controls, and SQL injection vulnerabilities.
Another critical capability in application vulnerability assessment practices includes package monitoring. Package monitoring is the practice of tracking and monitoring software packages or dependencies used in an application. It

involves keeping a close watch on the versions, vulnerabilities, and updates of these packages to ensure the security and integrity of the application.
Modern applications often rely on third-party libraries, frameworks, or modules that provide pre-built functionalities. These packages can introduce vulnerabilities if they contain outdated or insecure components. Package monitoring helps identify such vulnerabilities and allows developers to take appropriate actions to mitigate potential risks.
That's it for this lesson. In this lesson, we discussed vulnerability scanning. We looked at the three stages of network vulnerability scanning. We also looked at credentialed and non-credentialed scans. We talked about application vulnerability scanning and package monitoring—methods used to enhance the overall security posture of an organization's software.

7.2.2 Conduct Vulnerability Scans (Demo Video)

Transcript:

Throughout this course, we talk about many scanning methods. Typically, you scan for a specific purpose, often to look for vulnerabilities. Vulnerability scanning attempts to find points that can be exploited on a computer or network.

A vulnerability scan can detect and sometimes classify system weaknesses in computers, networks, and other equipment. The scan can also predict the effectiveness of countermeasures that you might implement. A scan may be performed by an IT department or someone hired to perform a penetration test. There are several good tools that perform vulnerability scanning. But if you're on a budget, you can use nmap with a script, the Nmap Scripting Engine, to do a vulnerability scan.
I'm on a Kali Linux system. This system comes with a whole bunch of scripts, but the one we want isn't included. That's not a problem; we can go out to the web and find it. To get the script, we'll go to Github.
I've already located the web page with the script. This is the one that we want. I'll open that link. Down here, in the description, it says that the script uses some well-known service to provide information about vulnerabilities. That's really not a great description, but I happen to know that it works pretty well.
To get the script, you need to click on this Clone or Download link and copy the address. Now, we'll open the terminal, type in 'git clone', paste in the rest of the address, and press Enter on the keyboard. I've already done that, so I won't do it again.
When I was doing this, I wasn't paying attention, and it went into the root directory, so I had to go in and move it to the /usr/share/nmap/scripts directory so nmap knows where to find it by default. If you have some basic Linux skills, it shouldn't be a problem to move to the directory before downloading the script.
All right. On my network, I have a Windows 10 system, and I've installed XAMPP on it with some web pages, php, forms, and so on. I want to put it out there facing the public internet. But I'm not sure if it has any vulnerabilities. This is my example. I could also have other systems that I can scan. But for this demo, I'm just going to scan this particular system. Since I'm familiar with this system, I already know the IP address. To begin, I want to just do a regular service scan with nmap to see what we get. For that, I'll type in 'nmap -sV 10.10.10.195' and press Enter. It takes a few seconds to get the results. When it's done, you can see that I have a few open ports and services running. As I already mentioned, this system is running a web server, Apache, along with other services, so I'm concerned about it. Now let's run this scan with the script.
To run the script, we'll type 'nmap - -script nmap-vulners -sV 10.10.10.195' and press Enter. This will take a few seconds. The scan is done. Now we have some additional information. Right here, we have some information under Apache and port 443. I'm going to hold down the Ctrl key and click on this hyperlink. It'll open Firefox, and I'll go to the page where I can read more about this particular vulnerability. I'll scroll down. Under the description, it tells me exactly what this issue is. It looks like someone could execute some code, so it's probably not something I want to let loose on the web until I do more research and fix any problems.
While we're in Firefox, let's jump back over to the Github tab. I'll go back a page. There are other scripts here that do some additional scans. If you have a test environment (and you should), you can download others and see what type of results you can get. Remember that you can run multiple scripts in one scan.
That's it for this demo. In this demo, we used nmap and the Nmap Scripting Engine to do a vulnerability scan on a Windows 10 system running XAMPP. First, we went to Github to locate the scripts we needed. Then we did a regular service scan with nmap. We did another scan using the script, and then we used the web to learn more about one of the vulnerabilities that the script found.

7.2.3 Scanning a Network with Nessus (Demo Video)

Transcript:

In this demonstration, we're going to work with the Nessus Vulnerability Scanner. Nessus is a very powerful security tool that you can use to scan for security vulnerabilities on your network.

Nessus is accessed via a web page on port 8834. I've installed Nessus on this machine, so I'm just connecting to localhost. For this demo, I'm running an evaluation copy of Nessus, but that's sufficient to show what it can do.
Let's start by logging in. After I log in, we see the home page. Here, we can see the various scans that have already been created. There are many different features and options in Nessus, but for this demo, I'll stick to showing how to create and run a network scan.
While I could create a custom scan profile, Nessus includes quite a few useful templates that can accomplish most of what I typically need. To create a scan, I'll start by selecting New Scan from the top right. Here, we can see some of the templates I mentioned. I'm going to run a basic network scan, so I'll select that template. This template scans a specified host or network range for known vulnerabilities.
To set up the scan, I'll need to provide a name for it, as well as a network address that I'd like to scan. I'm going to name this 'home gateway scan' since I'll be scanning the computers at my house with it. My gateway is located at 10.0.0.1, so that's what I'll supply as the target. There are many other options that I could configure, including port discovery and authenticated scanning, but I just want to do a simple scan today. Since everything I need is configured, I'll click Save.
Now that I've returned to the home page, I can see that the new scan has been created, but isn't running yet. I can click Launch to start it.
The scan's title changes to bold, and a spinning green icon appears to show that the scan is active. I can watch the scan live by clicking on it. This screen shows information about the different hosts that are discovered and the vulnerabilities that are found. Nessus also keeps a historical log of all the times a scan has been run so that I can compare how vulnerabilities have improved on my network over time.
This scan is going to take a while, so I'll pause the demo until it's finished.
Now that the scan has finished, we can see that my gateway has quite a few vulnerabilities. Luckily, most of them are labelled as informational. This means that the tool found information that could be useful for a security administrator to know about a host, but that the information doesn't immediately indicate a problem. An example of the type of data that would be labelled as informational is a list of open network ports. Network ports have to be open on some machines for them to function properly, but opening unexpected ports can cause problems.
More serious vulnerabilities are color-coded from green to red. Low-severity vulnerabilities aren't as interesting as high or critical vulnerabilities, so I'll take a look at one of the worst issues Nessus found.
When I select a vulnerability, it shows me information about the issue. In many cases, software creators patch high and critical vulnerabilities, and a relevant CVE entry is listed. Nessus provide a description and a suggested solution for many issues, as well as relevant information related to the vulnerability.
And that's it for this demo. We discussed how to conduct a vulnerability scan with Nessus. We looked at what Nessus is and what it does. And we ran a vulnerability scan on my home network and explored the results.

7.2.4 Scanning a Network with OpenVAS (Demo Video)

Transcript:

Greenbone's OpenVAS is a powerful network security scanner equipped with various tools, including an intuitive graphical user interface. At its core, OpenVAS comprises a server packed with a suite of network vulnerability tests designed to identify security issues within remote systems and applications. This versatile tool serves offensive and defensive security professionals, enabling them to assess and uncover potential attack surfaces. In this demonstration, we'll continue to use the traditional name OpenVAS when referring to the scanner.

Before diving into this demonstration, I downloaded OpenVAS as a virtual appliance and imported it into a hypervisor, setting it up as a virtual machine. The system is configured, and I've run some tests, so we're ready to proceed. I'm at the login screen and will use the credentials created during the setup phase.

After logging in, I'm directed to the Dashboard screen. This provides a quick overview of previous scans and their status. It's important to note that security scanners, like antivirus software, rely on up-to-date definitions. Let's examine this briefly.
I'll navigate to the Administration tab and click on Feed Status. This column shows that all our feeds are up to date. I'm particularly interested in the CVEs (Common Vulnerabilities and Exposures).
Now, let's return to the Dashboards. Please remember that this is a test system, and the abundance of red indicates severe issues. In a typical network, such extensive red flags would be highly unusual and cause serious concern. Now, let's explore the Scans menu.
Clicking on Tasks reveals a list of all the tasks I've created. For instance, the "Scan All" task scans my entire subnet. In the Reports section, you'll see that I've run this task three times, each resulting in a report. You can also check the last time the task ran and its severity level. If I wish to run the task, I can start by clicking the appropriate arrow icon. Additional actions include deleting, editing, cloning, or exporting the task.
Now, let's create a new task. There are a couple of ways to do this. We can use the Task Wizard to quickly create a task with default settings, requiring only the entry of the IP address or hostname, whether a single target or an entire range of devices. Another approach is to click on New Task.
I'll name the task Scan All since it's intended for scanning the entire network. Under Scan Targets, we specify what we want to scan. I've created a new target by clicking here.
This opens the New Target configuration window. I'll name it Scan All and enter the network address and subnet I want to scan. Here, you can also specify exclusions if there are devices on your network you wish to skip. Credentials can be provided for a more comprehensive scan.
After clicking Save, we return to the New Task window. You'll see that the Scan All target is already selected. For this task, we'll leave the rest of the settings at their default values and click Save.
Now, looking at the list of tasks, you can find the Scan All task. Since it's new, there have yet to be any reports. While I'm ready to run this task, I want to explore other menu items first.
Under the Scans menu, you'll find options like Reports, Results, Vulnerabilities, and more. Reports display a list of accumulated reports since installing OpenVAS, while Results reveal vulnerabilities discovered during scans. For example, you can identify an operating system that has reached its end of life and the affected host system. Addressing these vulnerabilities is crucial, as outdated systems are unlikely to receive automatic patches.
Returning to the Scans menu, we'll examine the Tasks. On the first page of the list, I'll select a specific report, focusing on a task that scanned IP addresses ending with "100," including a vulnerable host called Metasploitable. I'll click on the report and head to the Results tab to view the issues detected by OpenVAS. These issues range from passwordless logins to an end-of-life operating system and a VNC brute force login vulnerability. Expanding on a particular point provides detailed information and suggests solutions.
Under the Scan menu, there's an extensive list of menu items, including Assets, Configuration, and more. Assets allow you to view scanned hosts, while Configuration lets you create targets and credentials. Lastly, the Administration and Help sections provide access to administration tasks and helpful resources.
This concludes our demonstration. We explored Greenbone's Open Vulnerability Assessment Scanner, also known simply as OpenVAS.

7.2.5 Vulnerability Scanning Facts

This lesson covers the following topics:

- Vulnerability scanning
- Application vulnerability scanning
- Package monitoring

Vulnerability Scanning

Vulnerability management is a cornerstone of modern cybersecurity practices aimed at identifying, classifying, remediating, and mitigating vulnerabilities within a system or network. A vulnerability scanner looks for weaknesses such as:

- Open ports.
- Active IP addresses.

- Running applications or services.
- Missing critical patches.
- Default user accounts that have not been disabled.
- Default or blank passwords.
- Misconfigurations.
- Missing security controls.

One crucial aspect of vulnerability management is vulnerability scanning, a systematic process of probing a system or network using specialized software tools to detect security weaknesses. Vulnerability scans are performed internally and externally to inventory vulnerabilities from different network viewpoints. Vulnerabilities identified during scanning are then classified and prioritized for remediation by security operations teams.

Vulnerability scanning also supports application security, as it helps to locate and identify misconfigurations and missing patches in software. Advanced vulnerability scanning techniques focused on application security include specialized application scanners, pen-testing frameworks, and static and dynamic code testing.

Vulnerability scanning tools like openVAS and Nessus offer a broad range of features to analyze network equipment, operating systems, databases, patch compliance, configuration, etc. While these tools are very effective, application security analysis warrants more specialized approaches. Several specialized tools exist to more deeply analyze how applications are designed to operate and can locate vulnerabilities not typically identified using generalized scanning approaches.

There are different options when running a vulnerability scan. The table below explains each of these options:

Vulnerability Scan Option	Description
Intrusive	An intrusive scan finds a potential vulnerability and then actively attempts to exploit it. This leads to more accurate results but cannot be done on a live system.
Non-intrusive	A non-intrusive scan is the more common type of scan performed. This method scans the network and lists all potential vulnerabilities but cannot validate if the system is vulnerable. This type of scan can be performed on live systems and requires the network defender to take additional actions.
Credentialed	A credentialed scan is given a user account with login rights to various hosts, plus whatever other permissions are appropriate for the testing routines. This sort of test allows much more in-depth analysis, especially in detecting when applications or security settings may be misconfigured. It shows what an insider attack, or an attack with a compromised user account, may achieve. A credentialed scan is a more intrusive type of scan than a non-credentialed scanning.
Non-credentialed	A non-credentialed scan proceeds by directing test packets at a host without being logged on to the OS or application. The view is the one the host exposes to an unprivileged user on the network. The test routines may be able to include things such as using default passwords for service accounts and device management interfaces, but they are not given privileged access. While you may discover more weaknesses with a credentialed scan, you will sometimes want to narrow your focus to that of an attacker who does not have specific high-level permissions or total administrative access. Non-credentialed scanning is the most appropriate technique for external assessment of the network perimeter or when performing web application scanning.

Application Vulnerability Scanning

Similarly, application vulnerability scanning describes a specialized method for identifying software application weaknesses. This includes static analysis (reviewing application code without executing it) and dynamic analysis (testing running

applications), which can identify issues like unvalidated inputs, broken access controls, and SQL injection vulnerabilities. Application vulnerability scanning is typically handled separately from general vulnerability scanning due to the unique nature of software applications and the specific types of vulnerabilities they introduce.

General vulnerability scanning is designed to detect system-wide or network-wide weaknesses, such as out-of-date software or misconfigured firewalls. In contrast, application vulnerability scanning evaluates the coding and behavior of individual software applications. It looks for issues like cross-site scripting (XSS), SQL injection, and insecure direct object references unique to software applications. These application-specific vulnerabilities require specialized tools and techniques to identify and mitigate and are generally different from those used in general vulnerability scanning. Applications frequently have their own release and update cycles, separate from the rest of the environment, necessitating a more targeted vulnerability management process.

Package Monitoring

Another critical capability in application vulnerability assessment practices includes package monitoring. Package monitoring is associated with vulnerability identification because it tracks and assesses the security of third-party software packages, libraries, and dependencies used within an organization to ensure that they are up-to-date and free from known vulnerabilities that malicious actors could exploit. Package monitoring is associated with managing software bill of materials (SBOM) and software supply chain risk management practices.

In an enterprise setting, package monitoring is typically achieved through automated tools and governance policies. Automated software composition analysis (SCA) tools track and monitor the software packages, libraries, and dependencies used in an organization's codebase. These tools can automatically identify outdated packages or packages with known vulnerabilities and suggest updates or replacements. They work by continuously comparing the organization's software inventory against various databases of known vulnerabilities, such as the National Vulnerability Database (NVD) or vendor-specific advisories.

In addition to these tools, organizations often implement governance policies around software usage. These policies may require regular audits of software packages, approval processes for adding new packages or libraries, and procedures for updating or patching software when vulnerabilities are identified.

7.2.6 Scan for Cleartext Vulnerabilities (Simulation)

Scenario

One of the content developers on the Engineering team uses an Embedthis GoAhead webserver for several devices that are maintained by their team.

In this lab, you need to scan the test machine where the Engineering team prepares deployments and complete the following tasks:

- Login to the CompTIA Vulnerability Scanner in Chrome.
  - URL: **http://192.168.0.52**
  - Username: **securityadmin**
  - Password: **P@ssw0rd**
- Use the CompTIA Vulnerability Scanner to scan the test machine found at 192.168.0.45.
- Answer the questions about any vulnerabilities found.

Explanation

Complete this lab as follows:

1. Access the CompTIA Vulnerability Scanner.
    a. URL: http://192.168.0.52
    b. Username: securityadmin
    c. Password: P@ssw0rd (with a zero, not the letter o)
    d. Click **Sign In** .
2. Scan the host at 192.168.0.45.
    a. Click on the **Targets** tab.
    b. Click on **Add Target** .
    c. Name the target **Webserver** or a name of your choice.
    d. In the **Hosts** field, enter **192.168.0.45** , then select **OK** .
    e. Click on the **Tasks** tab.
    f. Click on **Add Task** .
    g. Enter **Scan Webserver** (or a name of your choice) for the **Name** .
    h. Select **Webserver** (or the name you chose) from the **Add Target** list box. Click **OK**
    i. Click the **Run** button to the right to start the scan.
3. View the Vulnerability Scanner report.
    a. Click the **Reports** tab.
    b. Review the contents of the report under **Webserver** (or the name you chose) and answer the questions by clicking the **Answer Questions** button at the top right.
    c. Click **Score Lab** .

7.2.7 Scan for FTP Vulnerabilities (Simulation)

Scenario

A server is used to transfer company financial data to remote branches using the FTP protocol. Since the data is sensitive to the company, you have been asked to scan the host for vulnerabilities.

In this lab, your task is to complete the following:

- Login to the CompTIA Vulnerability Scanner in Chrome.
    - URL: **http://192.168.0.52**
    - Username: **securityadmin**
    - Password: **P@ssw0rd**
- Using the CompTIA Vulnerability Scanner, scan the server found at 192.168.0.46.
- Answer the questions presented about what the Vulnerability Scanner finds.

Explanation

Complete this lab as follows:

1. Access the CompTIA Vulnerability Scanner.
    a. URL: http://192.168.0.52
    b. Username: securityadmin
    c. Password: P@ssw0rd (with a zero, not the letter o)
    d. Click **Sign In** .
2. Scan the host at 192.168.0.46.
    a. Click on the **Targets** tab.
    b. Click on **Add Target** .
    c. Name the target **FTP** or a name of your choice.
    d. In the **Hosts** field, enter **192.168.0.46** , then select **OK** .
    e. Click on the **Tasks** tab.
    f. Click on **Add Task** .
    g. Enter **FTP scan** (or a name of your choice) for the **Name** .

h. Select **FTP** (or the name you chose) from the **Add Target** list box. Click **OK**
i. Click the **Run** button to the right to start the scan.

3. View the Vulnerability Scanner report.
    a. Click the **Reports** tab.
    b. Review the contents of the report under **FTP** (or the name you chose) and answer the questions by clicking the **Answer Questions** button at the top right.
    c. Click **Score Lab** .

7.2.8 Scan for TLS Vulnerabilities (Simulation)

Scenario

An older server has been providing file sharing for Windows, Linux, and MacOS clients to the Sales team.

In this lab, you need to scan the file server to ensure it is secure by completing the following tasks:

- Login to the CompTIA Vulnerability Scanner in Chrome.
    - URL: **http://192.168.0.52**
    - Username: **securityadmin**
    - Password: **P@ssw0rd**
- Use the CompTIA Vulnerability Scanner to scan the test machine found at 192.168.0.46.
- Answer the questions about any vulnerabilities found.

Explanation

Complete this lab as follows:

1. Access the CompTIA Vulnerability Scanner.
    a. URL: http://192.168.0.52
    b. Username: securityadmin
    c. Password: P@ssw0rd (with a zero, not the letter o)
    d. Click **Sign In** .
2. Scan the host at 192.168.0.46.
    a. Click on the **Targets** tab.
    b. Click on **Add Target** .
    c. Name the target **File Server** or a name of your choice.
    d. In the **Hosts** field, enter **192.168.0.46** , then select **OK** .
    e. Click on the **Tasks** tab.
    f. Click on **Add Task** .
    g. Enter **File Server** (or the name you chose) for the **Name** .
    h. Select **File Server** (or the name you chose) from the **Add Target** list box. Click **OK**
    i. Click the **Run** button to the right to start the scan.
3. View the Vulnerability Scanner report.
    a. Click the **Reports** tab.
    b. Review the contents of the report under **File Server** (or the name you chose) and answer the questions by clicking the **Answer Questions** button at the top right.
    c. Click **Score Lab** .

7.2.9 Scan for Windows Vulnerabilities (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You are performing vulnerability scans on your network. Mary is the primary administrator for the network and the only person authorized to perform local administrative actions. The company network security policy requires complex passwords for all users. It is also required that Windows Firewall is enabled on all workstations. Sharing personal files is not allowed.

In this lab, your task is to:

- Login to the CompTIA Vulnerability Scanner in Chrome.
  - URL: http://192.168.0.52
  - Username: securityadmin
  - Password: P@ssw0rd
  - Select **Sign In**
- Create a target for the Office2 workstation (192.168.0.34).
- Create a task and run a vulnerability scan for the Office2 workstation.
- View the report for the scan task you created.
- Remediate the vulnerabilities found in the report for Office2. Use Computer Management, Settings, and File Explorer to make needed changes.
- Re-run a vulnerability scan to make sure all of the issues are resolved.

Explanation

Complete this lab as follows.

1. Login to the CompTIA Vulnerability Scanner on ITAdmin in Chrome.
   - URL: http://192.168.0.52
   - Username: securityadmin
   - Password: P@ssw0rd
   - Select **Sign In**
2. Create a Target for Office2 (192.168.0.34)
   a. Select **Targets** , and then **Add Target** .
   b. Name: Office2 (or a name of your choice)
   c. Hosts: 192.168.0.34
   d. Select **OK**
3. Create a Task to scan the Office2 target and run the task
   a. Select **Tasks** , and then **Add Task** .
   b. Name: **Scan Office2** (or a name of your choice)
   c. To the right of Add Target, select **Office2** (or whatever name you chose)
   d. Select **OK** .
   e. Select **Run** .
4. View the Report for the Task
   a. Select Reports and view the report for Office2 (or whatever name you chose).
   b. Review the results to determine which issues you need to resolve on Office2.
5. Access local users using Office2's Computer Management console.
   a. From the top left, select **Floor 1** .
   b. Under Office 2, select **Office2** .
   c. From Office2, right-click **Start** and select **Computer Management** .
   d. Expand and select **Local Users and Groups** > **Users** .
6. Rename the Administrator user account.
   a. Right-click **Administrator** and select **Rename** .
   b. Enter a ***new name*** of your choice and press **Enter** .
7. Disable the Guest account.
   a. Right-click **Guest** and select **Properties** .
   b. Select **Account is disabled** , and then select **OK** .
8. Set a new password for Mary.
   a. Right-click **Mary** and select **Set Password** .

b. Select **Proceed** .
c. Enter a new **password** of your choice (12 characters or more).
d. Confirm the new **password** , and then select **OK** .
e. Select **OK** .

Ideally, you would create a policy that requires passwords with 12 characters or more with special characters and mixed cases.

9. Configure Mary's password to expire and to change at next logon.
    a. Right-click **Mary** and select **Properties** .
    b. Clear **Password never expires** .
    c. Select **User must change password at next logon** , and then select **OK** .
10. Unlock Susan's account and remove her from the Administrators group.
    a. Right-click **Susan** and select **Properties** .
    b. Clear **Account is locked out** , and then select **Apply** .
    c. Select the **Member of** tab.
    d. Select **Administrators** .
    e. Select **Remove** .
    f. Select **OK** .
    g. Close Computer Management.
11. Enable Windows Firewall for all profiles.
    a. Right-click **Start** , and then select **Settings** .
    b. Select **Network & Internet** .
    c. From the right pane, scroll down and select **Windows Firewall** .
    d. Under Domain network, select **Turn on** .
    e. Under Private network, select **Turn on** .
    f. Under Public network, select **Turn on** .
    g. Close all open Windows.
12. Remove a file share.
    a. From the taskbar, select **File Explorer** .
    b. From the left pane, select **This PC** .
    c. From the right pane, double-click **Local Disk (C:)** .
    d. Right-click **MyMusic** and select **Properties** .
    e. Select the **Sharing** tab.
    f. Select **Advanced Sharing** .
    g. Clear **Share this folder** .
    h. Select **OK** .
    i. Select **Close** .
13. Use the CompTIA Vulnerability Scanner from ITAdmin to verify that all of the issues on the Office2 computer were resolved.
    a. From the top left, select **Floor 1** .
    b. Under IT Administration, select **ITAdmin** .
    c. In the CompTIA Vulnerability Scanner, under Tasks, select **Rerun** .
    d. Select Reports and view the report for Office2 (or whatever name you chose)
    e. Review the results to determine if you need to resolve additional issues on Office2.

7.2.10 Scan for Linux Vulnerabilities (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You need to use a vulnerability scanner to check for security issues on your Linux computers.

In this lab, your task is to:

- Login to the CompTIA Vulnerability Scanner in Chrome.
    - URL: **http://192.168.0.52**
    - Username: **securityadmin**
    - Password: **P@ssw0rd**
- Create a target for the Linux computers on IP range **192.168.0.60 - 192.168.0.69**
- Answer the first question
- Create a task and run a vulnerability scan for the Linux range.
- View the report for the scan task you created.
- Answer the remaining questions.

Explanation

Complete this lab as follows:

1. Login to the CompTIA Vulnerability Scanner on ITAdmin in Chrome.
    - URL: **http://192.168.0.52**
    - Username: **securityadmin**
    - Password: **P@ssw0rd**
    - Select **Sign In**
2. Create a Target for theIP address range of 192.168.0.60 through 192.168.0.69.
    a. Select **Targets** , and then **Add Target** .
    b. Enter a Name of your choice.
    c. For Hosts, enter **192.168.0.60 - 192.168.0.69**
    d. Select **OK** .
    e. Expand the target you created to view the detected hosts.
    f. Answer Question 1.
3. Create a Task to scan the Linux range target and run the task
    a. Select **Tasks** , and then **Add Task** .
    b. Enter a Name of your choice
    c. To the right of Add Target, select your target.
    d. Select **OK** .
    e. Select **Run** .
4. View the Report for the Task
    a. Select Reports and view the report for your scan task
    b. Review the results and then answer Questions 2 through 5.

7.2.11 Scan for Domain Controller Vulnerabilities (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You are performing vulnerability scans on your network. Use the CompTIA Vulnerability Scanner tool to run a vulnerability scan on the CorpDC domain controller.

In this lab, your task is to:

- Login to the CompTIA Vulnerability Scanner in Chrome.
    - URL: http://192.168.0.52
    - Username: securityadmin
    - Password: P@ssw0rd
- Create a target for the CorpDC server (192.168.0.11).
- Create a task and run a vulnerability scan for the CorpDC server.
- View the report for the scan task you created.
- Remediate the vulnerabilities in the Default Domain Policy using Group Policy Management on CorpDC.

- Re-run a vulnerability scan to make sure all of the issues are resolved.

Explanation

While completing this lab, use the following information:

Area	Policy	Setting
Password Policy	Enforce password history	**24 Passwords**
	Minimum password age	**1 Day**
	Minimum password length	**14 Characters**
Account Lockout Policy	Reset account lockout counter after	**60 Minutes**
Event Log	Retention method for application log	**Do not overwrite events (clear log manually)**
	Retention method for security log	**Do not overwrite events (clear log manually)**
	Retention method for system log	**Do not overwrite events (clear log manually)**
System Services	DCOM Server Process Launcher	**Disabled**
	Task Scheduler	**Disabled**

Complete this lab as follows:

1. Login to the CompTIA Vulnerability Scanner on ITAdmin in Chrome.
   - URL: http://192.168.0.52
   - Username: securityadmin
   - Password: P@ssw0rd
2. Create a Target for CorpDC (192.168.0.11)
   a. Select **Targets** , and then **Add Target** .
   b. Name: **CorpDC** (or a name of your choice)
   c. Hosts: **192.168.0.11**
   d. Select **OK**
3. Create a Task to scan the CorpDC (or the name you chose) target and run the task
   a. Select **Tasks** , and then **Add Task** .
   b. Name: **Scan CorpDC** (or a name of your choice)
   c. Next to Add Target, select **CorpDC** (or the name you chose)
   d. Select **OK** .
   e. Select **Run** .
4. View the Report for the Task
   a. Select Reports and view the report for CorpDC (or the name you chose)
   b. Review the results to determine which issues you need to resolve on CorpDC.

5. Access the CorpDC server.
    a. From the top left, select **Floor 1** .
    b. Under Networking Closet, select **CorpDC** .

    If you need to return to the ITAdmin computer to review the report:

    1. From the top left, select **Floor 1** .
    2. Under IT Administration , select **ITAdmin** .

6. Access and edit the **CorpNet.local Default Domain Policy** .
    a. From Server Manager, select **Tools** > **Group Policy Management** .
    b. Maximize the window for better viewing.
    c. Expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** .
    d. Right-click **Default Domain Policy** and then select **Edit** .
    e. Maximize the window for better viewing.
7. Remediate the password policy issues in **Account Policies** .
    a. Under Computer Configuration, expand **Policies** > **Windows Settings** > **Security Settings** > **Account Policies** .
    b. From the left pane, select **Password Policy** .
    c. From the right pane, double-click the ***policy*** .
    d. Select **Define this policy setting** .
    e. Enter the ***password setting*** and then select **OK** .
    f. Repeat steps 7c-7e for each additional password policy.
8. Remediate the reset account lockout counter issue in **Account Policies** .
    a. From the left pane, select **Account Lockout Policy** .
    b. From the right pane, double-click **Reset account lockout counter after** .
    c. Select **Define this policy setting** .
    d. Enter **60** minutes and then select **OK** .
9. Remediate the Event Log issues.
    a. From the left pane, select **Event Log** .
    b. From the right pane, double-click the ***policy*** .
    c. Select **Define this policy setting** .
    d. Select **Do not overwrite events (clear log manually)** and then select **OK** .
    e. Repeat steps 9b–9d for each additional Event Log policy.
10. Remediate System Services issues.
    a. From the left pane, select **System Services** .
    b. From the right pane, double-click the ***policy*** .
    c. Select **Define this policy setting** .
    d. Make sure **Disabled** is selected, and then select **OK** .
    e. Repeat steps 10b-10d for the remaining System Services policy.
11. Use the CompTIA Vulnerability Scanner from ITAdmin to verify that all the issues on the CorpDC server were resolved.
    a. From the top left, select **Floor 1** .
    b. Under IT Administration, select **ITAdmin** .
    c. In the CompTIA Vulnerability Scanner, under Tasks, select **Rerun** .
    d. Select **Reports** and view the report for CorpDC.
    e. Review the results to determine if you need to resolve additional issues on CorpDC.

7.2.12 Practice Questions (Section Quiz)

q_vuln_scanning_credentialed_01_secp8

Which of the following are key purposes of running a credentialed scan in a vulnerability assessment? (Select two.)

Answers:

- ***Testing routines**
- Public network access
- ***Compromised user account**
- External network perimeter
- Unprivileged user access

Explanation:

The following answers are correct:

- Testing routines are a key aspect of a credentialed scan. The scan is given a user account with login rights to various hosts, plus whatever other permissions are appropriate for the testing routines. This allows for a more in-depth analysis of the system.
- A credentialed scan simulates what an insider attack, or an attack with a compromised user account, may achieve. This is because it is given the same level of access as these potential threats.

Public network access is not a key aspect of a credentialed scan. Credentialed scans are typically performed internally and do not require public network access.

The external network perimeter is not a key aspect of a credentialed scan. While the external network perimeter may be scanned during a vulnerability assessment, this is typically done using non-credentialed scans.

Unprivileged user access is not a key aspect of a credentialed scan. Credentialed scans are given privileged access to the system to allow for a more thorough analysis.

q_vuln_scanning_credentialed_02_secp8

You are a cybersecurity analyst at a large corporation. Your team has been tasked with conducting a vulnerability assessment of the company's internal network. You have been given the option to perform either a credentialed or non-credentialed scan.

Which of the following factors would most strongly suggest that a credentialed scan is the appropriate choice for this situation?

Answers:

- The company's network has recently been targeted by a series of external cyber attacks.
- The company has a large number of third-party applications installed on its network.
- ***The company has recently implemented a new security policy that restricts the use of administrative privileges.**
- The company's IT department has recently installed a new patch management system.

Explanation:

A credentialed scan is designed to provide a more in-depth analysis of the network, including detecting misconfigurations in security settings. If the company has recently implemented a new security policy that restricts the use of administrative privileges, a credentialed scan would be able to assess the impact of this policy and identify any potential vulnerabilities that may have been introduced.

While external cyber attacks are a serious concern, they are not the primary reason to choose a credentialed scan. A non-credentialed scan would be more appropriate for assessing the vulnerabilities visible to an external attacker.

While third-party applications can introduce vulnerabilities, a credentialed scan is not specifically designed to assess these. Application vulnerability scanning would be more appropriate for this purpose.

While a new patch management system could potentially introduce vulnerabilities, a credentialed scan is not specifically designed to assess these. A patch management audit would be more appropriate for this purpose.

q_vuln_scanning_dynamic_analysis_secp8

An application security analyst at a software company is assessing a new software application before releasing it to customers. Before deciding on the BEST approach for the assessment, the analyst recalls that there are different methods of analysis to evaluate the software's security posture.

The analyst wants to assess the software's running state to identify potential vulnerabilities during its execution.

Considering the preference to evaluate the software in its running state and identifying vulnerabilities during execution, which type of examination should the analyst primarily rely on?

Answers:

- Static code review
- Manual penetration testing
- ***Dynamic analysis**
- Source code fingerprinting

Explanation:

Dynamic analysis evaluates the software application in its running state and looks for vulnerabilities during its execution, which aligns with the analyst's requirement in the scenario.

Static code review evaluates the software's source code, bytecode, or application binaries without executing the software. While it is a valuable method, static code review does not meet the analyst's preference to assess the software while running.

Manual penetration testing involves actively probing for vulnerabilities in a running application, but is broader than just analyzing the software's execution and can involve various techniques not limited to the software's runtime behavior.

Source code fingerprinting involves snippets within files that match content in source files found in third-party components. Source code fingerprints act as identifiers of likely third-party content within the scanned file. While a valid approach, it is not the best approach for this scenario.

q_vuln_scanning_noncredentialed_01_secp8

Which of the following are key areas of focus for a non-credentialed scan in a vulnerability assessment? (Select two.)

Answers:

- ***External network perimeter**
- Privileged user access
- ***Unprivileged user access**
- Internal network access
- Compromised user account

Explanation:

The following answers are correct:

- The external network perimeter is a key focus of a non-credentialed scan. Non-credentialed scans are often used to assess the security of the network perimeter from an external viewpoint, simulating the perspective of an attacker who does not have specific high-level permissions or total administrative access.
- Unprivileged user access is a key focus of a non-credentialed scan. Non-credentialed scans simulate the view that the host exposes to an unprivileged user on the network.

Privileged user access is not a key focus of a non-credentialed scan. Non-credentialed scans are performed without logging into the system or application, thus they do not have privileged user access.

Internal network access is not a key focus of a non-credentialed scan. While non-credentialed scans can be used internally, they are often used to assess the security of the network perimeter from an external viewpoint.

A compromised user account is not a key focus of a non-credentialed scan. Non-credentialed scans are performed without logging into the system or application, thus they do not simulate an attack from a compromised user account.

q_vuln_scanning_noncredentialed_02_secp8

You are a cybersecurity analyst at a financial institution. Your team has been tasked with conducting a vulnerability assessment of the company's external network perimeter.

You have been given the option to perform either a credentialed or non-credentialed scan.

Which of the following factors would MOST strongly suggest that a non-credentialed scan is the appropriate choice for this situation?

Answers:

- ***The company's network has recently been targeted by a series of external cyber attacks.**
- The company has a large number of third-party applications installed on its network.
- The company has recently implemented a new security policy that restricts the use of administrative privileges.
- The company's IT department has recently installed a new patch management system.

Explanation:

A non-credentialed scan is designed to mimic the perspective of an external attacker who does not have privileged access to the network. If the company's network has recently been targeted by external cyber attacks, a non-credentialed scan would be the most appropriate choice to identify the vulnerabilities that an external attacker could potentially exploit.

While third-party applications can introduce vulnerabilities, a non-credentialed scan is not specifically designed to assess these. Application vulnerability scanning would be more appropriate for this purpose.

A non-credentialed scan does not require administrative privileges and therefore would not be affected by a policy that restricts the use of these privileges. A credentialed scan, which provides a more in-depth analysis of the network, would be more appropriate in this situation.

While a new patch management system could potentially introduce vulnerabilities, a non-credentialed scan is not specifically designed to assess these. A patch management audit would be more appropriate for this purpose.

q_vuln_scanning_nvd_secp8

As a cybersecurity analyst, you are tasked with identifying known vulnerabilities in the third-party software packages, libraries, and dependencies used within your organization.

Which of the following would be the MOST effective tool for accomplishing this task?

Answers:

- Software Bill of Materials (SBOM)
- Software composition analysis (SCA)
- ***National Vulnerability Database (NVD)**
- Intrusion detection system (IDS)

Explanation:

The National Vulnerability Database (NVD) is a U.S. government repository of standards-based vulnerability management data. It includes databases of security checklist references, security-related software flaws, misconfigurations, product names, and impact metrics. It is the most effective tool for identifying known vulnerabilities in third-party software packages, libraries, and dependencies.

A Software Bill of Materials (SBOM) is a comprehensive list of components in a piece of software. It provides detailed information about each component, including its source and version number. While an SBOM is useful for understanding what components are in your software, it does not provide information about known vulnerabilities.

Software composition analysis (SCA) tools are used to identify open source components in code and check them against known vulnerability databases. While SCA is a valuable tool in a cybersecurity toolkit, it does not provide a comprehensive list of known vulnerabilities that exist in third-party software packages, libraries, and dependencies.

An intrusion detection system (IDS) is designed to monitor network traffic for suspicious activity and issue alerts when such activity is detected. While an IDS is an important part of a comprehensive cybersecurity strategy, it is not specifically designed to identify known vulnerabilities in third-party software packages, libraries, and dependencies.

q_vuln_scanning_package_monitoring_secp8

You are a cybersecurity analyst at a large organization. You've noticed that several third-party software packages used within your organization have not been updated in a while.

What is the MOST appropriate action to take?

Answers:

- Ignore the issue as it's not your responsibility to update third-party software.
- Update the software packages immediately without informing anyone
- ***Inform your manager about the issue and suggest implementing automated package monitoring.**
- Delete the outdated software packages from the system.

Explanation:

Informing your manager about the issue and suggesting the implementation of automated package monitoring is the most appropriate action. Package monitoring is associated with vulnerability identification because it tracks and assesses the security of third-party software packages, libraries, and dependencies used within an organization to ensure that they are up-to-date and free from known vulnerabilities.

Ignoring the issue is not the right approach. Outdated software packages can pose a security risk as they might contain known vulnerabilities that malicious actors could exploit.

Updating the software packages immediately without informing anyone is not the right approach. It's important to communicate with your team and follow the proper procedures for updating software.

Deleting the outdated software packages from the system is not the right approach. These packages might be necessary for certain operations within the organization. Instead, they should be updated to the latest version.

q_vuln_scanning_sbom_secp8

As a cybersecurity analyst, you are tasked with improving the security of your organization's software applications. One of your responsibilities is to ensure that all third-party software packages, libraries, and dependencies used within your organization are up-to-date and free from known vulnerabilities.

Which of the following would be the MOST effective tool for accomplishing this task?

Answers:

- Software composition analysis (SCA)
- ***Software Bill of Materials (SBOM)**
- National Vulnerability Database (NVD)
- Intrusion detection system (IDS)

Explanation:

A Software Bill of Materials (SBOM) is a comprehensive list of components in a piece of software. It provides detailed information about each component, including its source, version number, and any known vulnerabilities. This makes it an effective tool for tracking and assessing the security of third-party software packages, libraries, and dependencies.

Software composition analysis (SCA) tools are used to identify open source components in code and check them against known vulnerability databases. While SCA is a valuable tool in a cybersecurity toolkit, it does not provide the comprehensive, detailed list of all components in a piece of software that an SBOM does.

The National Vulnerability Database (NVD) is a U.S. government repository of standards-based vulnerability management data. While it is a valuable resource for identifying known vulnerabilities, it does not provide the comprehensive, detailed list of all components in a piece of software that an SBOM does.

An intrusion detection system (IDS) is designed to monitor network traffic for suspicious activity and issue alerts when such activity is detected. While an IDS is an important part of a comprehensive cybersecurity strategy, it is not specifically designed to track and assess the security of third-party software packages.

q_vuln_scanning_scanners_secp8

Which of the following statements about network vulnerability scanners is true?

Answers:

- Network vulnerability scanners, such as Tenable Nessus and OpenVAS, are designed to test only servers and switches.
- Network vulnerability scanners only identify vulnerabilities but do not suggest any remediation techniques.
- Network vulnerability scanners do not depend upon a database of known software and configuration vulnerabilities.
- ***Network vulnerability scanners can test common operating systems, desktop applications, and server applications.**

Explanation:

Network vulnerability scanners are configured with information about known vulnerabilities and configuration weaknesses for typical network hosts. These scanners can test common operating systems, desktop applications, and server applications. This is useful for general-purpose scanning, but some types of applications might need more rigorous analysis.

The following statements about network vulnerability scanners are NOT true:

- Network vulnerability scanners like Tenable Nessus and OpenVAS are designed to test a variety of network hosts, including client PCs, mobile devices, servers, routers, and switches, not just servers and switches.
- Network vulnerability scanners not only identify vulnerabilities but also suggest remediation techniques. They compile a report about each vulnerability in their database that was found to be present on each host, categorize each identified vulnerability, assign an impact warning, and most tools also suggest remediation techniques.
- Network vulnerability scanners do depend upon a database of known software and configuration vulnerabilities. They use this database to identify vulnerabilities present on each host during the scanning process.

q_vuln_scanning_true_statement_secp8

Which of the following statements about vulnerability scanning is true?

Answers:

- Vulnerability scanning is a process of identifying, classifying, and ignoring vulnerabilities within a system or network.
- Network vulnerability scanners, such as Tenable Nessus and OpenVAS, are designed to test only servers and switches.
- Non-credentialed scans are more intrusive and provide a more in-depth analysis than credentialed scans.
- ***Package monitoring is a critical capability in application vulnerability assessment practices as it tracks and assesses the security of third-party software packages, libraries, and dependencies.**

Explanation:

Package monitoring is indeed a critical capability in application vulnerability assessment practices. It tracks and assesses the security of third-party software packages, libraries, and dependencies used within an organization to ensure that they are up-to-date and free from known vulnerabilities that malicious actors could exploit.

The following statements are NOT true:

- Vulnerability scanning is a process of identifying, classifying, remediating, and mitigating vulnerabilities within a system or network, not ignoring them.
- Network vulnerability scanners like Tenable Nessus and OpenVAS are designed to test a variety of network hosts, including client PCs, mobile devices, servers, routers, and switches, not just servers and switches.
- Non-credentialed scans are less intrusive than credentialed scans. Non-credentialed scans proceed by directing test packets at a host without being logged on to the OS or application, providing a view that the host exposes to an unprivileged user on the network. On the other hand, credentialed scans are given a user account with login rights to various hosts, allowing for a more in-depth analysis.

scan_vuln_cleartext_cve

Which CVE was reported for the discovered vulnerability?

Answers:

- ***CVE-2019-16645**
- CVE-2020-19841
- CVE-2022-05040
- CVE-2021-21101

Explanation:

The only CVE listed is CVE-2019-16645

scan_vuln_cleartext_solution

Which of the following are suggested possible solutions to remediate the vulnerability? (Select three)

Answers:

- ***Upgrade to a newer release**
- ***Remove the product**
- ***Replace the product with another one**
- Install patch 2.5.6.235
- Migrate users to a different service

scan_vuln_ftp_num

How many vulnerabilities were found on the FTP server?

Answers:

- ***2**
- 1
- 3
- 4

scan_vuln_ftp_components

Which of the following services or networking components were identified as being vulnerable? (Select two)

Answers:

- ***The FTP service**
- ***ICMP packet responses**
- The Apache webserver
- The SMB service
- The OpenSSH service

scan_vuln_tls_cve

Which CVEs were reported for the discovered vulnerability? (Select two)

Answers:

- CVE-2019-16645
- ***CVE-2011-3389**
- ***CVE-2015-0204**
- CVE-2021-2101
- CVE-2023-4311

Explanation:

The only CVE listed is CVE-2019-16645

scan_vuln_tls_solution

Which of the following are suggested possible solutions to remediate the vulnerability?

Answers:

- ***Disable TLSv1.0 and TLSv1.1 protocols in favor of the TLSv1.2+ protocols.**
- Remove TLS/SSL from the system
- Upgrade all clients to Windows 10+
- Encrypt all email

scan_vuln_linux_01

Which Linux computers were discovered by the scanner on IP address range 192.168.0.60 - 192.168.0.69?

Answers:

- ***192.168.0.60**
- 192.168.0.61
- ***192.168.0.62**
- 192.168.0.63
- 192.168.0.64
- ***192.168.0.65**
- 192.168.0.66
- 192.168.0.67
- ***192.168.0.68**
- 192.168.0.69

scan_vuln_linux_02

Which vulnerabilities are present on all the computers in the range? (Select two)

Answers:

- ***ICMP Timestamp Reply Information Disclosure**
- ***TCP Timestamps Information Disclosure**
- Cleartext Transmission of Sensitive Information via HTTP
- Telnet Unencrypted Cleartext Login
- VNC Server Unencrypted Data Transmission

scan_vuln_linux_03

For the Linux computer with the 192.168.0.60 IP address, which vulnerabilities should be remediated immediately? (Select two)

Answers:

- ***rlogin Passwordless Login**
- ***Operating System (OS) End of Life (EOL) Detection**
- TCP Timestamps Information Disclosure
- ICMP Timestamp Reply Information Disclosure

scan_vuln_linux_05

What is the recommended solution for the Telnet Unencrypted Cleartext Login vulnerability?

Answers:

- ***Replace Telnet with SSH**
- Update Telnet to latest version
- Use TLS v1.2+
- Update the SSL/TLS certificate

scan_vuln_linux_04

Which vulnerability has no known solution and will likely not get fixed?

Answers:

- ***GoAhead Server HTTP Header Injection Vulnerability**
- Cleartext Transmission of Sensitive Information via HTTP
- SSL/TLS: Deprecated TLSv1.0 and TLSv1.1 Protocol Detection
- Telnet Unencrypted Cleartext Login

7.3 Alerting and Monitoring

As you study this section, answer the following questions:

- What is the importance of monitoring a network?
- What tools are available to monitor a network for vulnerabilities?
- What role does SIEM play in network security?
- What is the risk of false positive alerts and alarms?

In this section, you will learn to:

- Analyze network traffic with Netflow.

The key terms for this section include:

Term	Definition

Network monitors	Collects data about network infrastructure appliances, such as switches, access points, routers, firewalls. This is used to monitor load status for CPU/memory, state tables, disk capacity, fan speeds/temperature, network link utilization/error statistics, and so on.
Netflow	A Cisco-developed means of reporting network flow information to a structured database.
System monitors	A system monitor implements the same functionality as a network monitor for a computer host. Like switches and routers, server hosts can report health status using SNMP traps.
System logs	Logs function both as an audit trail of actions and (if monitored regularly) provide a warning of intrusion attempts. Log review is a critical part of security assurance.
Vulnerability scanners	A vulnerability scanner will report the total number of unmitigated vulnerabilities for each host. Consolidating these results can show the status of hosts across the whole network and highlight issues with a particular patch or configuration issue.
Antivirus	Antivirus software detects malware by signature regardless of type, though detection rates can vary quite widely from product to product.
Data loss prevention	Data loss prevention (DLP) mediates the copying of tagged data to restrict it to authorized media and services.
Security information and event management (SIEM)	Software designed to manage security data inputs and provide reporting and alerting. The core function of a SIEM tool is to collect and correlate data from network sensors and appliance/host/application logs.
Reporting	A managerial control that provides insight into the security system's status.
Alert tuning	Correlation rules that reduce the incidence of false positive alerts and alarms.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions.
	• Obfuscation
	  ○ Tokenization
	  ○ Data masking
	2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
	• Monitoring
	3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

- Infrastructure considerations
    - Sensors

3.3 Compare and contrast concepts and strategies to protect data.

- General data considerations
    - Data at rest
    - Data in transit
    - Data in use
- Methods to secure data
- Encryption
    - Masking
    - Tokenization
    - Permission restrictions

4.1 Given a scenario, apply common security techniques to computing resources.

- Monitoring

4.3 Explain various activities associated with vulnerability management.

- Identification methods
    - Vulnerability scan

4.4 Explain security alerting and monitoring concepts and tools.

- Monitoring computing resources
    - Systems
    - Infrastructure
- Activities
    - Log aggregation
    - Alerting
    - Reporting
    - Archiving
    - Alert response and remediation/validation
    - Alert tuning
- Tools
    - Security Content Automation Protocol (SCAP)
    - Benchmarks
    - Agents/agentless
    - Security information and event management (SIEM)
    - Antivirus
    - Data loss prevention (DLP)
    - Simple Network Management Protocol (SNMP) traps
    - NetFlow
    - Vulnerability scanners

4.9 Given a scenario, use data sources to support an investigation.

- Log data
    - Firewall logs

	• Endpoint logs • OS-specific security logs • IPS/IDS logs • Network logs • Data sources • Dashboards 5.1 Summarize elements of effective security governance. • Procedures • Playbooks 5.6 Given a scenario, implement security awareness practices. • Reporting and monitoring
TestOut Security Pro	3.1 Harden computer systems 3.1.2 Configure anti-virus protection 4.2 Implement Encryption Technologies 4.2.1 Encrypt data communications 4.2.2 Encrypt files 5.1 Implement logging and auditing 5.1.2 Enable device logs

7.3.1 Alerting and Monitoring (Lesson Video)

Transcript:

Alerting and monitoring play a critical role in cybersecurity. Continuously monitoring systems and networks allow organizations to detect potential threats and breaches early. These monitoring systems trigger timely alerts, enabling security teams to take immediate action to isolate affected systems and implement recovery plans to protect crucial data.

Alerting and monitoring tools are used to generate alerts in real time. These tools often come equipped with dashboards for visualizing data. They also have advanced analytics capabilities for deeper insights into network security. There's a wide range of alerting and monitoring tools available. Let's take a moment to look at just a few.
First are network monitoring tools. These tools scrutinize and manage the functionality of a network. They help identify network bottlenecks and often predict potential issues before they become serious problems.
Next is a flow collector. Flow collectors are a means for recording metadata and statistics about network traffic rather than recording each frame. NetFlow is a network protocol developed by Cisco for collecting IP traffic information and monitoring network traffic.
System monitors provide information about the resources and performance of a system. They can monitor CPU usage, disk activity, and memory usage. System logs record the detailed operations of a system or network. They're critical for understanding system activities and identifying abnormalities.
Application and Cloud Monitors supervise the operation of applications and cloud-based services. They help ensure optimal performance and uptime while also tracking unusual activity. Vulnerability Scanners scan a system or network to identify any security weaknesses. They're an integral part of maintaining a secure IT environment, as they provide the first line of defense by identifying vulnerabilities before they can be exploited.

Data Loss Prevention (DLP) tools detect and prevent data breaches, exfiltration, or unwanted destruction of sensitive data. They also ensure that end users don't send sensitive or critical information outside the corporate network. Lastly, a SIEM, or Security Information and Event Management, is a comprehensive solution that provides real-time analysis of security alerts generated by networks and applications.
One of the functions of a vulnerability scan is to assess the configuration of security controls and application settings and permissions compared to established benchmarks. This is known as benchmarking. The scanner might try to identify whether there's a lack of controls that might be considered necessary or any system misconfiguration that would make the controls less effective or ineffective, such as antivirus software not being updated or management passwords left configured to the default. This testing requires specific information about best practices in configuring the application or security control.
These best practices are provided by listing the controls and appropriate configuration settings in a template.
Security Content Automation Protocol (SCAP) allows compatible scanners to determine whether a computer meets a configuration baseline. SCAP uses several components to accomplish this function.
Open Vulnerability and Assessment Language, or OVAL, is an XML schema describing system security state and querying vulnerability reports and information. Extensible Configuration Checklist Description Format, or XCCDF, is an XML schema for developing and auditing best practice configuration checklists and rules. Previously, best practice guides might've been written in prose for systems administrators to apply manually. XCCDF provides a machine-readable format that can be applied and validated using compatible software.
That's it for this lesson. In this lesson, we talked about alerting and monitoring. We reviewed several tools that can help with this. We also discussed the importance of establishing benchmarks for effective monitoring and alerting.

7.3.2 Alerting and Monitoring Facts

This lesson covers the following topics:

- Alerting and monitoring
- Alerting and monitoring tools

Alerting and Monitoring

Alerting and monitoring play a critical role in cybersecurity. Continuously monitoring systems and networks allow organizations to detect potential threats and breaches early. These monitoring systems trigger timely alerts, enabling security teams to take immediate action to isolate affected systems and implement recovery plans to protect crucial data.

Alerting and Monitoring Tools

The following chart describes several different commonly used alerting and monitoring tools:

Tool	Description
Network monitors	Distinct from network traffic monitoring, a network monitor collects data about network infrastructure appliances, such as switches, access points, routers, firewalls. This is used to monitor load status for CPU/memory, state tables, disk capacity, fan speeds/temperature, network link utilization/error statistics, and so on. Another important function is a heartbeat message to indicate availability. This data might be collected using the Simple Network Management Protocol (SNMP). An SNMP trap informs the management system of a notable event, such as port failure, chassis overheating, power failure, or excessive CPU utilization. The threshold for triggering traps can be set for each value. This provides a mechanism for alerts and alarms for hardware issues. As well as supporting availability, network monitoring might reveal unusual conditions that could point to some kind of attack.

Tool	Description
	As distinct from network traffic monitoring, a network monitor collects data about network infrastructure appliances, such as switches, access points, routers, firewalls. This is used to monitor load status for CPU/memory, state tables, disk capacity, fan speeds/temperature, network link utilization/error statistics, and so on. Another important function is a heartbeat message to indicate availability. This data might be collected using the Simple Network Management Protocol (SNMP). An SNMP trap informs the management system of a notable event, such as port failure, chassis overheating, power failure, or excessive CPU utilization. The threshold for triggering traps can be set for each value. This provides a mechanism for alerts and alarms for hardware issues. As well as supporting availability, network monitoring might reveal unusual conditions that could point to some kind of attack.
Netflow	A flow collector is a means of recording metadata and statistics about network traffic rather than recording each frame. Network traffic and flow data may come from a wide variety of sources (or probes), such as switches, routers, firewalls, and web proxies. Flow analysis tools can provide features such as the following: • Highlighting of trends and patterns in traffic generated by applications, hosts, and ports. • Alerting based on detection of anomalies, flow analysis patterns, or custom triggers. • Visualization tools that show a map of network connections and make interpretation of patterns of traffic and flow data easier. • Identification of traffic patterns revealing rogue user behavior, malware in transit, tunneling, or applications exceeding their allocated bandwidth. • Identification of attempts by malware to contact a handler or command & control (C&C) channel. NetFlow is a Cisco-developed means of reporting network flow information to a structured database. NetFlow has been redeveloped as the IP Flow Information Export (IPFIX) IETF standard (tools.ietf.org/html/rfc7011). A particular traffic flow can be defined by packets sharing the same characteristics, referred to as keys. A selection of keys is called a flow label, while traffic matching a flow label is called a flow record. A flow label is defined by packets that share the same key characteristics, such as IP source and destination addresses and protocol type. These five bits of information are referred to as a 5-tuple. A 7-tuple adds the input interface and IP type of service data. Each exporter caches data for newly seen flows and sets a timer to determine flow expiration. When a flow expires or becomes inactive, the exporter transmits the data to a collector.
System monitors and logs	A system monitor implements the same functionality as a network monitor for a computer host. Like switches and routers, server hosts can report health status using SNMP traps. Logs are one of the most valuable sources of security information. A system log can be used to diagnose availability issues. A security log can record both authorized and unauthorized uses of a resource or privilege. Logs function both as an audit trail of actions and (if monitored regularly) provide a warning of intrusion attempts. Log review is a critical part of security assurance. Only referring to the logs following a major incident is missing the opportunity to identify threats and vulnerabilities early and to respond proactively. Logs typically associate an action with a particular user. This is one of the reasons why it is critical that users do not share login details. If a user account is compromised, there is no means of tying events in the log to the actual attacker.

Tool	Description
Application and cloud monitors	SNMP offers limited functionality. There are numerous proprietary monitoring solutions for infrastructure, application, database, and cloud environments. Some are designed for on-premises and some for cloud, while some support hybrid monitoring of both types of environments. An application monitor will include a basic heartbeat test to verify that it is responding. Other factors to monitor include number of sessions and requests, bandwidth consumption, CPU and memory utilization, and error or security alert conditions. Cloud monitors will assess different facets of cloud services, such as network bandwidth, virtual machine status, and application health.
Vulnerability scanners	A vulnerability scanner will report the total number of unmitigated vulnerabilities for each host. Consolidating these results can show the status of hosts across the whole network and highlight issues with a particular patch or configuration issue.
Antivirus	Most hosts should be running some type of antivirus scan (A-V) software. While the A-V moniker remains popular, these suites are better conceived of as endpoint protection platforms (EPPs) or next-gen A-V. These detect malware by signature regardless of type, though detection rates can vary quite widely from product to product. Many suites also integrate with user and entity behavior analytics (UEBA) and use AI-backed analysis to detect threat actor behavior that has bypassed malware signature matching. Antivirus will usually be configured to block a detected threat automatically. The software can be configured to generate a dashboard alert or log via integration with a SIEM.
Data loss prevention	Data loss prevention (DLP) mediates the copying of tagged data to restrict it to authorized media and services. As with antivirus scanning, monitoring statistics for DLP policy violations can show whether there are issues, especially where the results show trends over time.

7.3.3 SIEM and SOAR (Lesson Video)

Transcript:

To protect our networks, we usually implement multiple automated devices such as IDSs, IPSs, firewalls, and security configurations. The problem with implementing these systems is that the amount of generated data becomes overwhelming for a human being to sort through. To help with this, we implement SIEM and SOAR systems. In this lesson, I'll go over each of these systems and how they work to help protect networks.

Security information and event management tools, or SIEM tools, work by gathering all sorts of information from a network and putting it together in one central location. More than just aggregating data, SIEM systems can also actively read all this information and determine if there's an actual threat! Let's see how this works.
Everything starts with the collectors. Log collectors are responsible for gathering event logs from security appliances, host systems, and applications. We can also add sensors to the system in order to capture network packets or data inputs from all the disparate systems on our network. This data is sent to the event collectors, and the event collectors send it all to the SIEM.
SIEM software takes this data, reads and analyzes it, and separates it into different categories such as logon attempts, database entries, port scans, network congestion, and more. You can review the reports to help find any suspicious network activity.
If data exceeds the defined thresholds of normal network activity, the SIEM sends an alert to the security administrator, who then can investigate it and take care of the threat as needed.
The next generation of SIEM systems are taking things to the next level. By implementing artificial intelligence and machine learning, these new systems can analyze user behavior and sentiment to determine if a threat exists. This can be used to detect threats like spear phishing attacks and insider threats.
SIEM systems are great at helping network administrators filter data and improve security monitoring. But, any alert still requires manual intervention.

The acronym SOAR stands for Security Orchestration, Automation, and Response. SOAR systems also gather and analyze data, but these systems take it to the next level. SOAR is a solution stack of compatible software programs that allow an organization to collect security-threat data from multiple sources and respond to low-level security events without human assistance. Let's break this down and see how these systems work.

SOAR systems gather the same information as SIEM systems do, but they also gather data from multiple third-party tools. The SOAR system coordinates these tools, sensors, and collectors to work together to gather as much relevant data as possible. This is the orchestration piece of a SOAR system.

You can set up a SOAR system to automate tasks that are routine, tedious, and time consuming, such as looking for and deleting phishing emails. This is usually configured using checklists called playbooks or a series of conditional steps called runbooks. Automating these tasks frees up time for your security team to focus on more important things.

Finally, a SOAR system is able to automatically respond to threats. For example, if malware is discovered, a SOAR system can identify the threat and quarantine it instead of just sending an alert.

SOAR systems help to reduce a security operation team's workload by automating workflows and handling low-level tasks automatically. You should use both SIEM and SOAR systems for improved security. Your SOAR system can likely respond to the low-level threats that the SIEM system uncovers, and your security team can take responsibility for the higher-level issues.

That'll wrap things up for now. In this lesson, we went over SIEM and SOAR systems. SIEM programs work by gathering all sorts of data from sensors and collectors. Then this data is analyzed, and alerts are sent out for potential threats. A SOAR system takes this one step further and is able to respond to low-level threats itself. You can configure SOAR systems to automate basic tasks that take up your organization's valuable time. Using both systems can greatly improve your network security.

7.3.4 SIEM and SOAR Facts

This lesson covers the following topics:

- Security Information and Event Management (SIEM)
- Alerting and monitoring activities
- Alerting
- Reporting
- Archiving
- Alert tuning

Security Information and Event Management (SIEM)

Software designed to manage security data inputs and provide reporting and alerting is often described as security information and event management (SIEM). The core function of a SIEM tool is to collect and correlate data from network sensors and appliance/host/application logs. In addition to logs from Windows and Linux-based hosts, this could include switches, routers, firewalls, IDS sensors, packet sniffers, vulnerability scanners, malware scanners, and data loss prevention (DLP) systems.

Collection is how the SIEM ingests security event data from various sources. There are three main types of security data collection:

- **Agent-based** —this approach means installing an agent service on each host. As events occur on the host, logging data is filtered, aggregated, and normalized at the host and then sent to the SIEM server for analysis and storage. Collection from Windows/Linux/macOS computers will use agent-based collection. The agent must run as a process and could use 50–500 MB of RAM, depending on the amount of activity and processing it does.
- **Listener/collector** —rather than installing an agent, hosts can be configured to push log changes to the SIEM server. A process runs on the management server to parse and normalize each log/monitoring source. This method is often used to collect logs from switches, routers, and firewalls, as these are unlikely to support agents. Some variant of the Syslog protocol is typically used to forward logs from the appliance to the SIEM.

- **Sensor** —as well as log data, the SIEM might collect packet captures and traffic flow data from sniffers. A sniffer can record network data using either a switch's mirror port functionality or some tap on the network media.

As distinct from collection, log aggregation refers to normalizing data from different sources to be consistent and searchable. SIEM software features connectors or plug-ins to interpret (or parse) data from distinct types of systems and to account for differences between vendor implementations. Each agent, collector, or sensor data source will require its own parser to identify attributes and content that can be mapped to standard fields in the SIEM's reporting and analysis tools. Another essential function is normalizing date/time zone differences to a single timeline.

Alerting and Monitoring Activities

When data has been collected and aggregated, the SIEM can implement alerting, reporting, and archiving activities. Note that these activities can be performed manually or automated using discrete tools for each security appliance. The advantage of a SIEM is to consolidate the activities into a single management interface. This consolidated functionality, referred to as a "single pane of glass," refers to the enhanced visibility into a complex environment that such software offers.

Alerting

Correlation means interpreting the relationship between individual data points to diagnose incidents of significance to the security team. A SIEM correlation rule is a statement that matches certain conditions. These rules use logical expressions, such as AND and OR, and operators, such as == (matches), < (less than), > (greater than), and in (contains). For example, a single-user login failure is not a condition that should raise an alert. Multiple user login failures for the same account, taking place within one hour, are more likely to require investigation and are a candidate for detection by a correlation rule.

Error.LoginFailure > 3 AND LoginFailure.User AND Duration < 1 hour

As well as the correlation between indicators observed in the collected data, a SIEM will likely be configured with a threat intelligence feed. This means that data points observed in the collected network data can be associated with known threat actor indicators, such as IP addresses and domain names.

Each alert will be dealt with by the incident response processes of analysis, containment, eradication, and recovery. When used in conjunction with a SIEM, two steps in alert response and remediation deserve particular attention:

- Validation during the analysis process is how the analyst decides whether the alert is a true positive and needs to be treated as an incident. A false positive generates an alert, but no actual threat activity exists.
- Quarantine isolates the source of indicators, such as a network address, host computer, or file.

Alert response and remediation steps will often be guided by a playbook that assists the analyst with applying all incident response processes for a given scenario. One of the advantages of SIEM and advanced security orchestration, authorization, and reporting (SOAR) solutions is to automate validation and remediation fully or partially. For example, a quarantine action could be available as a mouse-click action via an integration with a firewall or endpoint protection product. Validation is made easier by correlating event data to known threat data and pivoting between sources, such as inspecting the packets that triggered a particular IDS alert.

Reporting

Reporting is a managerial control that provides insight into the security system's status. A SIEM can assist with reporting activity by exporting summary statistics and graphs. Report formats and contents are usually tailored to meet the needs of different audiences:

- Executive reports provide a high-level summary for decision-makers. This guides planning and investment activity.

- Manager reports provide cybersecurity and department leaders with detailed information. This guides day-to-day operational decision-making.
- Compliance reports provide whatever information is required by a regulator.

Determining which metrics are most useful for reporting is always very challenging. The following types illustrate some common use cases for reporting:

- Authentication data, such as failed login attempts and critical file audit data.
- Hosts with missing patches and/or configuration vulnerabilities.
- Privileged user account anomalies include out-of-hours use or excessive requests for elevated permissions.
- Trend reporting to show changes to key metrics over time.

A SIEM can be used for two types of reporting:

- Alerts and alarms detect the presence of threat indicators in the data and can be used to start incident cases. Day-to-day management of alert reporting forms a large part of an analyst's workload.
- Status reports communicate data about the level of threat or number of incidents being raised and the effectiveness of security controls and response procedures. This type of reporting can be used to inform management decisions. It might also be required for compliance reporting.

A SIEM will ship with several preconfigured dashboards and reports, but it will also make tools available for creating custom reports. It is critical to tailor the information presented in a dashboard or report to meet the needs and goals of its intended audience. If the report contains an overwhelming amount of data or irrelevant information, it will not be possible to identify remediation actions quickly.

Archiving

A SIEM can enact a retention policy to keep historical log and network traffic data for a defined period. This allows for retrospective incident and threat hunting and can be a valuable source of forensic evidence. It can also meet compliance requirements to hold archives of security information. SIEM performance will degrade if excessive data is kept available for live analysis. A log rotation scheme can be configured to move outdated information to archive storage.

Alert Tuning

Correlation rules are likely to assign a criticality level to each match. Examples include the following:

- Log only — an event is produced and added to the SIEM's database, but it is automatically classified.
- Alert — the event is listed on a dashboard or incident handling system for an agent to assess. The agent analyzes the event data and either dismisses it to the log or validates it and starts an incident case.
- Alarm — the event is automatically classified as critical, and a priority alarm is raised. This might mean emailing an incident handler or sending a text message.

Alert tuning is necessary to reduce the incidence of false positives. False positive alerts and alarms waste analysts' time and lower productivity. Alert fatigue is when analysts are so consumed with dismissing numerous low-priority alerts that they miss a single high-impact alert that could have prevented a data breach. Analysts can become more preoccupied with looking for a quick reason to dismiss an alert than with adequately evaluating the alert. Reducing false positives is difficult, however, firstly because there is not a simple dial to turn for overall sensitivity and secondly because reducing the number of rules that produce alerts increases the risk of false negatives.

There is also a concept of true negatives. This is a measure of events that the system has properly allowed. Metrics for false and true negatives can be used to assess the performance of the alerting system.

Some of the techniques used to manage alert tuning include the following:

- Refining detection rules and muting alert levels — If a particular rule generates multiple dashboard notifications, the rule's parameters can be adjusted to reduce this, perhaps by adding more correlation factors. Alternatively, the alert can be muted to log-only status or configured only to produce a single notification for every 10 or 100 events.
- Redirecting sudden alert "floods" to a dedicated group — Changes in the network can cause a rule to produce far more alerts than it should. Once confirmed that this is a false positive, rather than "spamming" each analyst's dashboard, it can be assigned to a dedicated agent or team to remediate.
- Redirecting infrastructure-related alerts to a dedicated group — Misconfigurations, such as deviance from a baseline, can cause continually high alert volumes. While these are important to fix, that is not the job of the incident response team and is better managed by an infrastructure team.
- Continuous monitoring of alert volume and analyst feedback — Managers should oversee the system and be aware of risks from alert fatigue. The experience of individual analysts can be utilized to reduce alert sensitivity, change the parameters of a given rule, or automate the processing of the rule using a SOAR solution.
- Deploying machine learning (ML) analysis — ML can rapidly analyze the data sets produced by SIEM. It can be used to monitor how analysts respond to alerts and attempt to automatically tune the ruleset to reduce false negatives without impacting true positives.

7.3.5 Analyze Network Traffic with Netflow (Demo Video)

Transcript:

NetFlow is a network protocol system originally created by Cisco. Netflow collects active IP network traffic as it flows in or out of an interface. The data is gathered and then analyzed to create a picture of network traffic flow and the volume of traffic.

So, why do we use Netflow? Well the NetFlow protocol is used by IT professionals as a network traffic analyzer to determine its point of origin, it's destination, the volume and the paths on the network. Before NetFlow's creation, network IT engineers and administrators would use Simple Network Management Protocol (SNMP) for network traffic analysis and monitoring.
Although Netflow is a feature was created by Cisco, there are open-source alternatives such as softflowd that can be installed to work with pfsense. This will allow us to capture network traffic and create a picture to better analyze in and out traffic.
In order to use softflowd we must install it from the package list. To do so we need to go to system then package manager. Tab over to Available packages. Just to make it easier so were not scrolling for a while type 'softflow' in the search bar and click search. Next, we can click the Install button and then confirm. When the software is installed successfully, we can then configure it.
To adjust our settings for softflowd we need to go to services, then softflowd. Make sure at the top it says enabled, Interface can either be LAN or WAN depending on what side of the firewall you would like to capture the data. It is possible to grab data from both if you would like however were just going to select WAN for now. The Host will be the NetFlow analyzer which happens to be '192.168.30.202' on our network. Port will the designated port of your choosing. Sometimes it is a good idea to use an alternative port for security purposes, but for now were using the standard '2055' port for NetFlow. We will leave the max flows to the default '8192'. The NetFlow version will depend on your analyzer's capability. If it doesn't support version ten then you may want to see what version, it does support. I'm going to flip this to version '9' because that is what our analyzer supports. Click save at the bottom.
Even though our settings are saved we want to make sure things are working right. There is a command line command you can run to show statistics on softflowd. Go to diagnostics then command prompt. In the Execute shell command field we can insert our command 'softflowctl -c /var/runsoftflowd.vmx0.ctl statistics'. The part that has vmx0 may differ if your network interface is labeled differently. Execute that command and now we can see some results. You may not a lot of traffic yet however you may want to check back later to verify flow packets and byte counts are going up. Now we are ready to send and configure NetFlow to our analyzer.
That is, it for this demo, in this demo we installed and configured softflowd.

7.3.6 Data Loss Prevention (Lesson Video)

Transcript:

Every business has sensitive data in its system, and protecting it is a high priority. A data leakage incident happens when sensitive data like credit card numbers, intellectual property, financial information, or proprietary company information is disclosed to an unauthorized person. In this lesson, we'll look at five approaches to data security, including data loss prevention, or DLP; masking; encryption; tokenization; and rights management.

Data can exist in one of three states, and it's important to use the right security approach for each state. The first state is while the data is in use on an endpoint system, like a workstation. The second state is while the data is in motion, such as when it's transmitted over the network or to the cloud. The third state is while the data rests on a storage medium, like a hard disk drive, or in a database.

Let's look at data protection through a DLP system. A DLP system works like a guard at the perimeter of your network, allowing unsensitive data to leave, but restricting sensitive data from being transmitted out of the company. A DLP analyzes the network traffic in accordance with the organization's security policies.

For example, an e-commerce retailer could use a network DLP solution to monitor for files containing credit card numbers. If one of those files were being transmitted out of the company, the DLP software would flag it as a potential security problem.

Next, let's look at the masking approach. Masking works by replacing sensitive data with realistic fictional data. There are different types of masking. We'll look at dynamic data masking, and then we'll look at static data masking.

Let's start with dynamic. Dynamic masking replaces original information with a mask that mimics the original in form and function, making it useful for data that's in use or in processing. For example, someone's name would be replaced with another name, or credit card numbers would be replaced with a random number that contains the same number of characters.

This method can be used to control which users can see the actual data. A bank could have third-party analytics performed on their accounts while masking the account numbers and clients' names. With dynamic data masking, the original data can be retrieved.

Another type of masking is static data masking. This type is helpful for data at rest in a database and can be specified by field or column. You may want to use this method if you're making copies of a database for testing, development, or reporting. The complex algorithms in static masking make the original data irretrievable through reverse-engineering, so making a masked copy may be a better choice than masking the original database.

All right. Next, let's briefly review encryption. Encryption happens when an algorithm changes plain text data into unreadable ciphertext. The encryption algorithm has a variable that's called a key. The authorized user that receives the encrypted data can decrypt it through the cipher key. This helps to protect data in motion.

Now let's look at a tokenization approach. Tokenization is similar to encryption and masking--it replaces actual data with something else. But tokenization uses a randomly generated alphanumeric character set called a token to replace the original data. The token server stores the original data and is protected by security measures like authentication and authorization protocols so that the original information is disclosed only when the correct token is presented. This method is frequently used for credit card numbers, bank accounts, medical records, and other personally identifiable information.

For example, when you have a credit card number stored on your mobile phone through an app, the app connects with the remote token server, which creates the token and replaces the credit card number stored in your phone.

Then, when you go to use your phone's app to pay for your purchase at the store, the store's point of sale terminal will contact a merchant acquirer. The merchant acquirer presents the token value to the remote token server. The server uses the token to map back to your actual credit card number and authorize the purchase. An authorization is sent to the merchant, and a message from the server is sent to your phone.

Finally, we have rights management. Rights management is data protection at the file level. With rights management, you identify sensitive files in the file system and embed them with your organization's security policy. The key benefit of this approach is that the security policy travels with the specific file even if it's moved or copied. You can continue to control access to the file, such as restricting who it can be transferred to, even when the file is no longer on your system.

Rights management falls into two categories: Digital Rights Management, or DRM, and Information Rights Management, or IRM.

DRM is file level management applied to rich media like music, videos, and software. This strategy uses security technologies such as encryption, permissions, product keys, limited install applications, and persistent online authentication to prevent editing, sharing, and unauthorized copying. The purpose is to protect copyrighted media and

software. For example, when a consumer purchases a software program, the program is not accessible without a product key provided by the manufacturer at the time of purchase. An example of continuous online authentication is when a consumer logs in to an online application or streams the information through an account that requires authentication.
Now let's move on to IRM, or Information Rights Management. It's also called Enterprise Rights Management sometimes. It focuses on business-to-business transfers for files such as documents, emails, spreadsheets, and financial data. Information rights management utilizes encryption and permissions to create rules for the files, which can allow or deny copying and pasting, editing, forwarding, and printing. An example is a contract document that only the recipient can open and digitally sign and is denied forwarding abilities.
That's it for this lesson. In this video, we review five approaches to data protection. First, we looked at data loss prevention. Next, we discussed two types of masking, dynamic and static. We talked about encryption, which uses a cipher with a key to encode information that can be decoded with a key by the receiver. Next, we discussed tokenization, which uses an alphanumeric value as a token to replace sensitive information that's protected by the token server. And finally, we went over rights management, which protects data through permissions at the file level that stay with the file even if it leaves your network.

7.3.7 DLP Facts

This lesson covers the following topics:

- Data loss prevention (DLP)

Smaller organizations might classify and type data manually to apply data guardianship policies and procedures. However, an organization that creates and collects large amounts of personal data usually needs automated tools to assist with this task. Protecting valuable intellectual property (IP) data may also be required. Data loss prevention (DLP) products automate the discovery and classification of data types and enforce rules so that data is not viewed or transferred without proper authorization. Such solutions will usually consist of the following components:

- Policy server — to configure classification, confidentiality, and privacy rules and policies, log incidents, and compile reports.
- Endpoint agents — to enforce policy on client computers, even when they are not connected to the network.
- Network agents — to scan communications at network borders and interface with web and messaging servers to enforce policy.

DLP agents scan content in structured formats, such as a database with a formal access control model, or unstructured formats, such as email or word processing documents. A file cracking process is applied to unstructured data to render it in a consistent scannable format. The transfer of content to removable media, such as USB devices, by email, instant messaging, or even social media, can be blocked if it does not conform to a predefined policy. Most DLP solutions can extend the protection mechanisms to cloud storage services, using either a proxy to mediate access or the cloud service provider's API to perform scanning and policy enforcement.

Remediation is the action the DLP software takes when it detects a policy violation. The following remediation mechanisms are typical:

- Alert only — copying is allowed, but the management system records an incident and may alert an administrator.
- Block — the user is prevented from copying the original file but retains access. The user may or may not be alerted to the policy violation, but it will be logged as an incident by the management engine.
- Quarantine — access to the original file is denied to the user (or possibly any user). This might be accomplished by encrypting the file or moving it to a quarantine area in the file system.
- Tombstone — the original file is quarantined and replaced with one describing the policy violation and how the user can re-release it.

When configured to protect a communications channel such as email, DLP remediation might take place using client-side or server-side mechanisms. For example, some DLP solutions prevent attaching files to the email before sending it. Others might scan the email attachments and message content, strip out specific data, or stop the email from reaching its destination.

7.3.8 Practice Questions (Section Quiz)

q_alert_monitoring_antivirus_secp8

A software technician is upgrading the alerting and monitoring tools for the organization's infrastructure and wants to employ new software to assist in detecting and removing infectious files and other forms of malware.

What type of software can the technician purchase to achieve this protection?

Answers:

- ***Antivirus (A-V)**
- Security content automation protocol
- Simple Network Management Protocol (SNMP) trap
- Data loss prevention

Explanation:

Antivirus (A-V) software can detect and remove infections and (in most cases) other types of malware, such as worms, Trojans, rootkits, adware, spyware, password crackers, network mappers, and denial-of-service (DoS) tools.

Security content automation protocol (SCAP) allows compatible scanners to determine whether a computer meets a configuration baseline.

A Simple Network Management Protocol (SNMP) trap informs the management system of a notable event, such as port failure, chassis overheating, power failure, or excessive central processing unit (CPU) utilization.

Data loss prevention (DLP) mediates the copying of tagged data to restrict it to authorized media and services and monitors statistics for policy violations.

q_alert_monitoring_application_secp8

Which tool assesses different facets of cloud services, such as network bandwidth, virtual machine status, and program health in a network environment?

Answers:

- Vulnerability scanner
- System monitor
- ***Application monitor**
- Data loss prevention (DLP) tool

Explanation:

An application monitor assesses an application's health, performance, and functionality, ensuring its smooth operation and detecting any potential issues.

A vulnerability scanner identifies and reports unmitigated vulnerabilities on hosts, aiding in the assessment of system security status. However, its primary focus is not evaluating application health or cloud services.

A system monitor is a hardware or software component used to monitor system resources and performance in a computer system. However, its primary focus is not evaluating application health or cloud services.

DLP tools constantly monitor and analyze data to identify potential violations of security policies and, if appropriate, stop them from continuing. However, its primary focus is not evaluating application health or cloud services.

q_alert_monitoring_continuous_secp8

A manufacturing company's security manager plans to implement corrective operational controls to mitigate potential security threats.

Which of the following instances would be the appropriate control?

Answers:

- ***Enabling continuous monitoring to disable abnormal accounts.**
- Regular penetration testing to uncover potential vulnerabilities.
- A firewall that prevents unauthorized access to the network.
- A security camera system monitoring the premises.

Explanation:

Enabling continuous monitoring to disable abnormal accounts is a corrective operational control. When detecting abnormal behavior, this control disables the account to prevent unauthorized access.

Penetration testing is more of a detective control than a corrective one. It identifies vulnerabilities but does not correct them directly.

A firewall is primarily a preventive control, not a corrective one. Its main function is to stop unauthorized access before it happens rather than correcting issues after they occur.

Security cameras are typically a deterrent and detective type of physical control, not an operational one. They can deter potential intruders and detect security incidents, but they do not correct issues directly.

q_alert_monitoring_flow_secp8

A company's IT department has noticed irregularities in network usage and resource allocation.

Which tool would be MOST beneficial in identifying patterns in network traffic, detecting anomalies, and visualizing network connections?

Answers:

- ***Flow collector**
- Network monitor
- SNMP trap
- Heartbeat message

Explanation:

Flow collectors record metadata and statistics about network traffic, thereby identifying trends and patterns, detecting anomalies, and providing visualization tools that simplify the interpretation of traffic data.

While network monitors can provide valuable information on the state of network appliances, they primarily focus on aspects like CPU/memory load status, disk capacity, network link utilization, and similar data.

Simple Network Management Protocol (SNMP) traps inform a management system of notable events, such as port failures or excessive CPU utilization, primarily dealing with hardware issues rather than traffic pattern analysis.

A heartbeat message indicates availability and does not directly analyze or interpret network traffic.

q_alert_monitoring_log_files_01_secp8

A digital forensics analyst at a healthcare company is investigating a case involving a potential internal data breach. The breach has led to unauthorized access and potential exposure of sensitive patient information.

The company uses a Security Information and Event Management (SIEM) tool that aggregates and correlates data from multiple sources. The analyst's task is to identify potential insider threats that could be responsible for the breach.

Given the nature of the breach, which combination of data sources should the analyst primarily consider for their investigation?

Answers:

- ***Log files generated by the OS components of client and server host computers, logs generated by applications and services running on hosts, and endpoint logs.**
- Packet captures system memory metadata and logs generated by endpoint security software installed on hosts.
- Log files generated by network appliances like switches, routers, and firewalls; application logs; and automated reports from the SIEM tool.
- Network traffic captured by sensors, logs generated by network-based vulnerability scanners, and firewall logs.

Explanation:

Client and server host operating system (OS) logs reveal system-level activities, while application and endpoint logs provide application usage and end-user activity insights. This combination identifies potential insider threats and provides a comprehensive view.

Packet captures and system memory metadata provides network and system activity data. Endpoint security logs reveal end-user threats.

Network appliance logs provide network activity insights, while application logs reveal unauthorized access or usage anomalies. System Information and Event Management (SIEM) tool reports summarize incidents but may miss system-level and endpoint activities.

Sensors capture detailed network activity views. Vulnerability scanner logs reveal exploitable vulnerabilities, while firewall logs provide traffic insights. However, this may need to include system-level and application activities.

q_alert_monitoring_log_files_02_secp8

A security operations analyst at a financial institution analyzes an incident involving unauthorized transactions. The analyst suspects that a malware infection on one of the endpoints might have led to the unauthorized access.

To identify the root cause and trace the activities of the suspected malware, which combination of data sources should the analyst primarily consider?

Answers:

- ***Endpoint logs, log files generated by the OS components of the affected host computer, and logs from the host-based intrusion detection system.**
- Network logs, packet captures, and logs generated by network-based vulnerability scanners.
- Firewall logs, system memory metadata, and automated reports from the SIEM tool.
- Logs from applications involved in the transactions, logs generated by the host's antivirus software, and /var/log/auth.log for authentication and authorization data.

Explanation:

Endpoint logs and operating system (OS) component log files comprehensively view potential malware activities on the affected end-user system.

Network logs, packet captures, and vulnerability scanner logs provide a network-level view but might not cover all system-level and application activities.

Firewall logs, system memory metadata, and System Information and Event Management (SIEM) reports provide network and system insights but might not cover all application-level activities.

Application logs and antivirus logs give insights into application usage and threats but might not apply to all hosts and could potentially miss network activities.

q_alert_monitoring_netflow_01_secp8

A network administrator at a large tech company has the task of enhancing the visibility into network traffic patterns in a distributed enterprise network.

The administrator wants to implement a solution that captures metadata and statistics about network traffic without recording each frame, with the goal of improving the company's security measures.

Which tool should the administrator consider implementing?

Answers:

- ***A NetFlow collector**
- A vulnerability scanner
- A simple network management protocol (SNMP) trap
- A data loss prevention

Explanation:

A NetFlow collector is the best solution. The NetFlow protocol collects and records metadata and statistics about network traffic, providing administrators with insights into traffic patterns.

A vulnerability scanner identifies security weaknesses in the network but does not capture metadata and statistics about network traffic.

A Simple Network Management Protocol (SNMP) trap alerts administrators of notable events in network devices in real time but does not capture metadata and statistics about network traffic.

A data loss prevention (DLP) system prevents data breaches by detecting potential data exfiltration transmissions but does not capture metadata and network traffic statistics.

q_alert_monitoring_netflow_02_secp8

A security analyst is monitoring a web application for potential security issues. After observing unusual behavior, the analyst wants to identify the source of the issue.

Which strategy would MOST effectively identify patterns revealing rogue user behavior, malware in transit, tunneling, or applications exceeding their allocated bandwidth?

Answers:

- ***Deploying a NetFlow collector to analyze network traffic related to the web application.**
- Implementing simple network management protocol (SNMP) traps in the web application.
- Using a vulnerability scanner to scan the web application.
- Installing an additional security information and event management (SIEM) system to handle the data from the web application.

Explanation:

The best strategy is to deploy a NetFlow collector to analyze network traffic related to the web application. NetFlow collectors capture metadata and statistics about network traffic to identify patterns and anomalies.

Implementing Simple Network Management Protocol (SNMP) traps in the web application is not the right approach. SNMP traps alert notable events in network devices but not patterns in web application traffic.

A vulnerability scanner can identify security weaknesses in the web application but may not effectively identify real-time traffic patterns related to the application.

Installing an additional SIEM system is unnecessary. A SIEM system can aggregate and analyze data from various sources, but a NetFlow collector can more effectively analyze network traffic patterns related to a specific application.

q_alert_monitoring_snmp_trap_secp8

After experiencing a catastrophic server failure in the headquarters building, what can the company use to monitor notable events such as port failure, chassis overheating, power failure, or excessive central processing unit (CPU) utilization?

Answers:

- ***Simple network management protocol (SNMP) trap**
- Data loss prevention
- Antivirus (A-V)
- Security content automation protocol

Explanation:

A Simple Network Management Protocol (SNMP) trap informs the management system of a notable event such as port failure, chassis overheating, power failure, or excessive central processing unit (CPU) utilization.

Data loss prevention (DLP) mediates the copying of tagged data to restrict it to authorized media and services and monitors statistics for policy violations.

Antivirus (A-V) software detects and removes infections and (in most cases) other types of malware, such as worms, Trojans, rootkits, adware, spyware, password crackers, network mappers, and denial of service (DoS) tools.

Security content automation protocol (SCAP) allows compatible scanners to determine whether a computer meets a configuration baseline.

q_siem_soar_alerts_and_alarms_secp8

As a cybersecurity analyst, you are responsible for monitoring the organization's security information and event management (SIEM) system. Your team has recently noticed an increase in suspicious network activity.

What type of SIEM reporting should you prioritize to effectively manage this situation?

Answers:

- Executive reports providing a high-level summary for decision-makers.
- Compliance reports providing information required by a regulator.
- ***Alerts and alarms detecting the presence of threat indicators in the data.**
- Status reports communicating data about the level of threat or number of incidents being raised.

Explanation:

Alerts and alarms detecting the presence of threat indicators in the data is the correct answer. In a situation where suspicious network activity is increasing, prioritizing alerts and alarms can help the team quickly identify and respond to potential threats. This type of reporting provides real-time information about potential security incidents, allowing for immediate action.

While executive reports provide a high-level summary for decision-makers and are important for informing strategic decisions, they may not provide the immediate, detailed information needed to address an increase in suspicious network activity.

Compliance reports providing information required by a regulator are crucial for maintaining regulatory compliance, but they may not provide the real-time, actionable information needed to respond to an immediate threat.

Status reports communicating data about the level of threat or number of incidents being raised can provide useful context about the overall security situation, but they may not provide the immediate, detailed information needed to respond to an increase in suspicious activity.

q_siem_soar_alert_response_secp8

A senior security analyst is refining the incident response processes for a large organization that recently implemented a security information and event management (SIEM) system. During a simulation of a cybersecurity incident, the analyst observed that the SIEM system generated several alerts that were false positives, leading to unnecessary consumption of resources.

On which step should the analyst focus to improve the efficiency of the alert response and remediation process?

Answers:

- ***Enhancing the validation and quarantine processes in the alert response**
- Increasing the number of correlation rules in the SIEM system

- Implementing additional threat intelligence feeds in the SIEM system
- Increasing the frequency of SIEM system reporting to capture more incidents

Explanation:

The analyst should focus on enhancing validation and quarantine processes in alert response. Validation reduces false positives, while quarantine isolates the source of indicators to manage potential incidents.

Simply increasing the number of correlation rules in the SIEM system is not the right solution, as it could lead to more false positives if not done judiciously.

Adding additional threat intelligence feeds in the SIEM system does not directly address false positives and could cause data overload and more false positives.

Increasing SIEM system reporting frequency does not directly improve alert response and remediation efficiency, as it does not address false positives in alerting.

q_siem_soar_collector_secp8

Which SIEM component is responsible for gathering all event logs from configured devices and securely sending them to the SIEM system?

Answers:

- Data handling
- SIEM alerts
- ***Collectors**
- Security automation

Explanation:

Collectors are responsible for gathering all event logs from configured devices and securely sending them to the security information and event management (SIEM) system. Collectors are basically the middleman between devices and the SIEM system.

The data handling component receives the data from the collectors and then reads, analyzes, and separates the data into different categories.

SIEM alerts are responsible for triggering alerts if any data exceeds the established thresholds.

Security automation is a feature of a SOAR system.

q_siem_soar_correlation_rule_secp8

A cybersecurity analyst uses a security information and event management (SIEM) tool to monitor network activity in a large organization.

During a shift, the analyst receives multiple alerts indicating the same user account is experiencing multiple login failures within the span of an hour.

Which of the following correlation rules likely triggered this alert?

Answers:

- Error.LoginFailure > 1 AND LoginFailure.User AND Duration < 1 day
- ***Error.LoginFailure > 3 AND LoginFailure.User AND Duration < 1 hour**
- Error.LoginFailure > 5 AND LoginFailure.User AND Duration < 30 minutes
- Error.LoginFailure > 2 AND LoginFailure.User AND Duration < 2 hours

Explanation:

Error.LoginFailure > 3 AND LoginFailure.User AND Duration < 1 hour - This correlation rule accurately represents the scenario described, which involves multiple login failures from the same user account within an hour.

Error.LoginFailure > 1 AND LoginFailure.User AND Duration < 1 day - While this correlation rule will flag repeated login failures, it is not as precise as the rule described in the scenario, specifically looking for more than three failures within one hour.

Error.LoginFailure > 5 AND LoginFailure.User AND Duration < 30 minutes - This correlation rule is stricter than the scenario described. It is looking for more than five login failures within 30 minutes, while the scenario involves multiple failures within an hour.

Error.LoginFailure > 2 AND LoginFailure.User AND Duration < 2 hours - This correlation rule is less precise than the one in the scenario. It looks for more than two failures within two hours, which is a broader timeframe than specified in the scenario.

q_siem_soar_detection_rules_secp8

A company tasks a cybersecurity manager with improving the efficiency of its security information and event management (SIEM) system.

The manager observes that the high number of false positive alerts causes alert fatigue among the analysts, potentially leading them to miss high-impact alerts.

Which combination of strategies should the manager consider implementing to tackle this issue effectively?

Answers:

- ***Refine detection rules, redirect sudden alert "floods" to a dedicated group, and continuously monitor alert volume and analyst feedback.**
- Assign all infrastructure-related alerts to the incident response team and increase the frequency of system reporting.
- Mute all alerts to log-only status and deploy additional threat intelligence feeds in the SIEM system.
- Increase the number of correlation rules and assign all alerts to a dedicated agent or team to remediate.

Explanation:

The best strategy is to refine detection rules, redirect sudden alert "floods" to a dedicated group, and continuously monitor alert volume and analyst feedback to address false positives and alert fatigue.

Assigning all infrastructure-related alerts to the incident response team is not the right approach. The infrastructure team should manage these alerts. Increasing system reporting frequency does not directly address false positives.

Muting all alerts to log-only status could lead to missed alerts. Deploying additional threat intelligence feeds does not directly address false positives.

Increasing the number of correlation rules does not necessarily reduce false positives and could increase them. Assigning all alerts to a dedicated agent or team does not specifically address false positives.

q_siem_soar_device_logs_secp8

A security administrator reviews the configuration of a newly implemented security information and event management (SIEM) system. The SIEM system collects and correlates data from various sources, such as network sensors, application logs, and host logs.

The administrator notices that some network devices, like switches and routers, do not directly support the installed agents for data collection.

What approach should the administrator consider to ensure the inclusion of these devices' logs in the SIEM system?

Answers:

- ***Configuring the devices to push log changes to the security information and event management (SIEM) server using a listener/collector approach**
- Implementing an additional data loss prevention (DLP) system for these devices
- Installing additional security information and event management (SIEM) servers to handle the data from these devices
- Running a vulnerability scanner on these devices to ensure they are compliant with the security information and event management (SIEM) system

Explanation:

Using a listener/collector approach to push log changes from switches, routers, and firewalls to the SIEM server is a good approach. A listener/collector approach includes logs from devices that do not support agents in the SIEM server.

An additional data loss prevention (DLP) system is not the appropriate solution for collecting logs from devices that do not support installed agents.

Installing additional SIEM servers is unnecessary since the problem relates to the method of log collection, not the capacity of the SIEM server.

Running a vulnerability scanner on these devices does not solve the log collection problem, as vulnerability scanners identify security weaknesses, not enable log collection for SIEM systems.

q_siem_soar_login_failures_secp8

An information security manager is fine-tuning a security information and event management (SIEM) system in a company that has recently experienced multiple cybersecurity incidents. The manager wants to ensure prompt detection of potential incidents for immediate investigation.

Which approach should the manager consider to optimize the system's alerting capability?

Answers:

- ***Configure the SIEM system to alert when multiple login failures for the same account occur within a specified time period.**
- Set the SIEM system to generate an alert for every single-user login failure.
- Enable the SIEM system to send an alert for every received threat intelligence feed.
- Arrange the SIEM system to archive all historical log data for retrospective incident and threat hunting.

Explanation:

The best approach is to configure the SIEM system to alert when multiple login failures for the same account occur within a specified time period, using correlation to interpret the relationship between data points and diagnose incidents for the security team.

Setting the SIEM system to generate an alert for every single-user login failure is not the right approach. This could lead to a high number of false positives.

Sending an alert for every received threat intelligence feed is not effective, as not all data may signify a potential incident.

Archiving is valuable for forensic analysis and compliance but does not address prompt detection of potential incidents.

q_siem_soar_logs_secp8

A security compliance manager at a large manufacturing company is investigating noncompliance incidents related to baseline security configurations across various hosts. The focus is identifying missing patches, understanding historical vulnerabilities, and assessing host configuration.

The manager needs to differentiate between remediated and unremediated vulnerabilities and correlate the findings with the date of the last scan.

What approach would BEST provide the required information for this investigation?

Answers:

- ***Utilize vulnerability scans and retrieving logs from a security event and incident management (SIEM).**
- Analyze network logs and application logs.
- Review packet captures and endpoint logs.
- Inspect operating system (OS)-specific security logs and automated reports.

Explanation:

Utilizing vulnerability scans and retrieving logs from a system information and event management (SIEM) would detect missing patches and noncompliance with baseline security configurations. The SIEM's ability to retrieve a list of logs for each host, with the date of the last scan, aligns with the investigation's objectives.

Generally, analyzing network and application logs doesn't provide specific details about missing patches and baseline security compliance.

Reviewing packet captures and endpoint logs offers details about network traffic and device activities but needs more depth for understanding historical vulnerabilities and baseline compliance.

Inspecting operating system (OS)-specific security logs and automated reports may provide general security information but need a detailed understanding of missing patches, noncompliance, and remediation status related to baseline configurations.

q_siem_soar_log_aggregation_secp8

A security analyst is optimizing a multinational company's security information and event management (SIEM) system. The system collects security event data from sources globally, and the analyst has noticed inconsistencies due to different time zones.

What should the analyst consider to ensure a consistent timeline across all logs for accurate event correlation?

Answers:

- ***Adjusting the log aggregation process in the SIEM system to normalize date/time zone differences.**
- Configuring the SIEM system to only collect data during the company's standard business hours.
- Implementing additional packet sniffers to collect network data uniformly.
- Installing agents on all data sources to ensure synchronization with the SIEM server's time zone.

Explanation:

Adjusting the log aggregation process in the SIEM system to normalize date/time zone differences is necessary. Normalizing date/time zone differences to a single timeline is an important function of the log aggregation process in a SIEM system.

Configuring the SIEM system to collect data only during business hours will not fix time zone inconsistencies and may cause the company to miss important security events.

Implementing additional packet sniffers to collect network data uniformly does not address the problem of time zone inconsistencies. Packet sniffers collect network data.

Installing agents on data sources for synchronization with the SIEM server's time zone does not fix the issue, as agents do not handle date/time zone normalization.

q_siem_soar_log_correlation_secp8

To optimize the enterprise security information and event management (SIEM) solution, a multinational 's chief information security officer (CISO) is strategizing. The SIEM system acquires data from diverse sources, including Linux and Windows servers, advanced switches, Next Generation Firewalls (NGFWs), and routers.

Which feature should the CISO prioritize improving in the SIEM solution to standardize the data and enhance its searchability?

Answers:

- ***Augmenting the log correlation mechanism in the SIEM solution.**
- Elevating the SIEM solution's threat-hunting capabilities.
- Upgrading the network-based data collection method in the SIEM solution.
- Integrating additional intrusion detection systems (IDS) into the network.

Explanation:

The most relevant improvement is augmenting the log correlation mechanism in the SIEM solution. Log correlation standardizes and makes data from various sources more searchable, directly addressing the CISO's objective.

Elevating the threat-hunting capabilities of the SIEM solution is crucial, but it does not directly influence the standardization and searchability of data from diverse sources.

Upgrading the network-based data collection method in the SIEM solution is not optimal. While this method assists in data collection, it does not directly address the issue of making data more standardized and searchable.

Integrating additional IDS into the network does not resolve the problem. While IDS collects network data, it does not directly contribute to standardization of data from diverse sources.

q_siem_soar_os_logs_secp8

A security operations center (SOC) manager notices a significant increase in unclassified events on the incident handler's security event and incident management (SIEM) dashboard. At the same time, someone or something raises the number of incidents.

The manager investigates these incidents further to ensure efficient and timely incident response.

Which combination of data sources would provide the MOST comprehensive view to support the manager's investigation?

Answers:

- ***OS-specific security logs, log files generated by applications and services running on hosts, and automated reports from the SIEM tool.**
- Firewall logs, network traffic captured by sensors, and log files generated by OS components of server host computers.
- Endpoint logs, automated reports from the SIEM tool, and metadata.
- Log files from network-based vulnerability scanners, application logs, and endpoint logs.

Explanation:

Operating system (OS)-specific security logs reveal system-level activities, while application and service logs on hosts provide application activity information. security information and event management (SIEM) tool reports summarize incidents, offering a balanced combination for comprehensive investigations.

Firewall logs and sensors provide detailed network activity views, while server-host OS logs offer server-side insights. However, this may not cover endpoint or application activities.

Endpoint logs reveal end-user system activities, and SIEM tool reports summarize incidents. However, metadata may need to provide more detail for thorough investigations.

Network-based vulnerability scanner logs show potential vulnerabilities, while application and endpoint logs reveal application and end-user system activities. However, this may not cover network-level activities.

q_siem_soar_playbook_secp8

Which of the following security orchestration, automation, and response (SOAR) system automation components is often used to document the processes and procedures that are to be used by a human during a manual intervention?

Answers:

- Runbook
- Orchestration
- Response
- ***Playbook**

Explanation:

Playbooks are linear checklists of required steps and actions that are to be taken to respond to an alert. While playbooks do support automated actions, they are often used to document the processes and procedures that are to be used by a human during a manual intervention.

Runbooks consist of a series of conditional steps to perform actions, such as sending notifications or threat containment. They are not used to document the processes and procedures that are to be used by a human during a manual intervention.

The orchestration component of the security orchestration, automation, and response (SOAR) system is responsible for gathering data and information from across the network. This is not used to document the processes and procedures that are to be used by a human during a manual intervention.

The response component of a SOAR system allows the system to automatically take actions against threats. It is not used to document the processes and procedures used by a human during a manual intervention.

q_siem_soar_retention_policy_secp8

As a cybersecurity analyst, you are tasked with managing the organization's security information and event management (SIEM) system. The system is experiencing performance degradation due to the volume of historical log and network traffic data.

What should you implement to manage the storage of this data effectively without compromising the system's performance?

Answers:

- Data compression techniques to reduce the size of the stored data.
- ***A retention policy to keep historical log and network traffic data for a defined period.**
- Increase the storage capacity of the SIEM system.
- Implement a data deletion policy to remove old data.

Explanation:

A retention policy to keep historical log and network traffic data for a defined period is the correct answer. A retention policy defines how long data should be kept based on its age, type, and relevance. This ensures that only necessary data is stored, improving system performance without compromising the ability to conduct retrospective incident and threat hunting.

While data compression techniques to reduce the size of the stored data can help to some extent, it does not address the root cause of the problem, which is the volume of data being stored. Additionally, compressed data may need to be decompressed for analysis, which can be time-consuming.

While increasing the storage capacity of the SIEM system might provide temporary relief, it does not address the underlying issue of managing the volume of data being stored. It could also lead to increased costs and further performance issues down the line.

While deleting old data can free up storage space, it may also remove valuable historical data that could be useful for future analysis or investigations. A retention policy is a more balanced approach that considers both storage limitations and the value of historical data.

q_siem_soar_security_logs_secp8

A technology firm's network security specialist notices a sudden increase in unidentified activities on the firm's security event and incident management (SIEM) incident tracking system.

An unknown entity or process also increases the number of reported incidents. The specialist decides to investigate these incidents.

Which combination of data sources would provide a balanced perspective to support the investigation?

Answers:

- ***System-specific security logs, which track system-level operations; logs generated by applications running on hosts; and real-time reports from the SIEM solution, summarizing incidents.**
- Gateway logs, which track incoming and outgoing network traffic; network interactions monitored by intrusion detection systems (IDS), which detect unauthorized activities; and logs from server OS components.
- User activity logs, which record user behaviors; daily summary reports from the SIEM solution; and high-level network overviews, providing a broad view of network activities.
- Logs from vulnerability assessment tools, which identify potential weaknesses; transaction logs from databases, tracking changes; and logs from mobile devices, recording device activities.

Explanation:

Security logs from specific systems and third-party applications provide insights into system and application operations. Real-time threat intelligence feeds from the SIEM solution to summarize incidents. Together, these sources offer a thorough view of in-depth investigations.

Gateway logs and intrusion detection system (IDS)-monitored interactions offer in-depth views of network activities, and virtualized server operating system (OS) logs provide server insights. However, this combination overlooks specific applications.

User activity logs focus on user behaviors, and daily summary reports and high-level network overviews need more granularity and depth for an exhaustive investigation.

Logs from host-based vulnerability assessment tools, database transaction logs, and mobile device logs would miss broader system and network-level activities required to understand the incidents.

q_siem_soar_soar_secp8

Which of the following systems is able to respond to low-level security events without human assistance?

Answers:

- SIEM
- IDS
- ***SOAR**
- Firewall

Explanation:

Security orchestration, automation, and response (SOAR) systems gather and analyze data like SIEM systems, but they take the analysis to the next level. SOAR is a solution stack of compatible software programs that allow an organization to collect data about security threats from multiple sources and respond to low-level security events without human assistance.

Security information and event management (SIEM) tools work by gathering different types of network information and data. This information is moved to one central place. SIEM systems are great tools that help network administrators filter data and improve security monitoring. Still, all alerts require manual intervention.

Intrusion detection systems (IDSs) can trigger alerts, but these systems do not respond to security threats on their own.

A firewall blocks traffic based on the configuration setup. However, firewalls do not respond to security threats on their own.

q_siem_soar_trend_reporting_secp8

You are a cybersecurity analyst working in a large organization that uses a security information and event management (SIEM) system.

The executive team has requested a report to help them understand the organization's security posture and make strategic decisions.

Which type of SIEM reporting metric would be MOST appropriate to use in this situation?

Answers:

- Authentication data, such as failed login attempts, and critical file audit data.
- Hosts with missing patches and/or configuration vulnerabilities.
- Privileged user account anomalies, such as out-of-hours use or excessive requests for elevated permissions.
- ***Trend reporting to show changes to key metrics over time.**

Explanation:

Trend reporting to show changes to key metrics over time is the correct answer. This type of reporting provides a high-level view of how the organization's security posture has evolved over time, which can help the executive team make strategic decisions. It can highlight patterns, show progress, and identify areas where further investment may be needed.

Authentication data, such as failed login attempts, and critical file audit data is important for operational security and immediate incident response, but it may not provide the strategic insight the executive team needs for long-term planning.

While hosts with missing patches and/or configuration vulnerabilities information is crucial for maintaining system security, it is more operational in nature and may not provide the high-level view required by the executive team.

Privileged user account anomalies, such as out-of-hours use or excessive requests for elevated permissions is a key metric for identifying potential insider threats or compromised accounts, but it is more focused on specific incidents rather than overall trends.

q_dlp_alert_only_secp8

You are the IT security manager at a mid-sized technology company. Your company uses a data loss prevention (DLP) system to protect sensitive data.

A senior developer, working on a critical project, attempts to copy a file containing non-confidential project code to a personal USB drive for working off-site. The DLP system detects this action as a violation of the company's data protection policy.

As the IT security manager, which DLP remediation action would you recommend to ensure the right balance between data protection and business continuity?

Answers:

- ***Alert only**
- Block
- Quarantine
- Tombstone

Explanation:

Alert only is the correct answer. This option would allow the copying of the file but would record the incident and possibly alert an administrator. Given that the file does not contain confidential information and the developer's role and need for off-site work, this option maintains business continuity while also keeping a record of the policy violation. This is the most appropriate solution in this context.

Block would prevent the user from copying the file while retaining their access to it. This could disrupt the developer's work and delay the critical project, which is not ideal given the non-confidential nature of the data.

Quarantine would deny the user access to the original file, possibly by encrypting it or moving it to a quarantine area in the file system. This is an excessive measure for non-confidential data and could significantly disrupt the developer's work and the project timeline.

Tombstone would quarantine the original file and replace it with a file describing the policy violation and how the user can re-release it. This is an extreme measure for a file that does not contain confidential data and could cause unnecessary delays and complications for the developer and the project.

q_dlp_dlp_secp8

Which of the following is a technology that tries to detect and stop sensitive data breaches, or data leakage incidents, in an organization?

Answers:

- ***Data loss prevention**
- Data transmission security
- Data hashing
- Public key cryptography

Explanation:

Data loss prevention (DLP) is a technology that tries to detect and stop sensitive data breaches, or data leakage incidents, in an organization. DLP is used to prevent sensitive data from being disclosed to an unauthorized person, whether it is deliberate or accidental.

Data transmission security is the use of secure protocols to encrypt data when it is transmitted.

Hashing takes a variable-length string (message) and compresses and transforms it into a fixed-length value. When received, a hash is decrypted into the actual output so the recipient can understand the message.

Public key cryptography infrastructure uses certificates, which are electronic documents that use a digital signature, to bind a public key with an identity.

q_dlp_end_dlp_secp8

Which of the following DLP implementations can be used to monitor and control access to physical devices on workstations or servers?

Answers:

- Network DLP
- ***Endpoint DLP**

- File-level DLP
- Cloud DLP

Explanation:

Endpoint data loss prevention (DLP) runs on end user workstations and servers. Endpoint DLP is also referred to as a Chinese Wall solution. This could be something as simple as restricting the use of USB devices. Many endpoint-based systems also provide application controls to prevent confidential information transmission and also provide some type of immediate feedback to the user. Giving feedback to the user is based on the concept that not all data leakage incidents are malicious. The employee might not realize that the security-policy violation is inappropriate. The intent is to deter the employee from a similar action in the future.

The following types of DLP are not designed to monitor and control access to physical devices on workstations or servers:

- Network DLP tracks and analyzes the organization's network activity and traffic, across a traditional network and the cloud; this includes monitoring e-mail, messaging and file transfers, to detect when business critical data is being sent in violation of the organization's information security policies.
- File-level DLP monitors, detects and blocks sensitive data from leaving an organization.
- Cloud DLP is designed to help you discover, classify, and protect your most sensitive data..

q_dlp_net_dlp_secp8

DLP can be implemented as a software or hardware solution that analyzes traffic in an attempt to detect sensitive data that is being transmitted in violation of an organization's security policies.

Which of the following DLP implementations analyzes traffic for data containing such things as financial documents, social security numbers, or key words used in proprietary intellectual property?

Answers:

- ***Network DLP**
- Endpoint DLP
- File-level DLP
- Cloud DLP

Explanation:

Network DLP is a software or hardware solution that is typically installed near the network perimeter. Network DLP analyzes network traffic in an attempt to detect sensitive data that is being transmitted in violation of an organization's security policies.

The following types of DLP are not designed to analyze network traffic in an attempt to detect sensitive data that is being transmitted in violation of an organization's security policies:

- Endpoint data loss prevention (DLP) runs on end user workstations and servers, not across the network.
- File-level DLP monitors, detects and blocks sensitive data from leaving an organization from individual servers or workstations.
- Cloud DLP is designed to help you discover, classify, and protect your most sensitive data.

q_dlp_tombstone_secp8

You are the IT security manager at a large corporation. Your company uses a data loss prevention (DLP) system to protect sensitive data.

An employee attempts to copy a file containing confidential client information to a personal USB drive. The DLP system detects this action as a violation of the company's data protection policy.

As the IT security manager, which DLP remediation action would you recommend to ensure the highest level of data protection while maintaining a record of the policy violation?

Answers:

- Alert only
- Block
- Quarantine
- ***Tombstone**

Explanation:

Tombstone is the correct answer. This option would quarantine the original file and replace it with a file describing the policy violation and how the user can re-release it. This not only prevents the user from accessing the sensitive data but also provides a clear record of the policy violation and a process for the user to regain access to the file. This is the most comprehensive solution and provides the highest level of data protection.

Alert only would allow the copying of the file but would record the incident and possibly alert an administrator. While this maintains a record of the policy violation, it does not prevent the data from being copied, which could lead to a data breach.

Block would prevent the user from copying the file while retaining their access to it. While this stops the immediate data breach, the user still has access to the file and could attempt to copy it again or find another way to extract the data.

Quarantine would deny the user access to the original file, possibly by encrypting it or moving it to a quarantine area in the file system. While this prevents the user from accessing the file, it does not provide a clear record of the policy violation.

7.4 Penetration Testing

As you study this section, answer the following questions:

- What is the purpose of a penetration test?
- What are the different types of penetration tests?
- What is the role of the purple team?
- Which document defines what is included in the penetration test?
- What is the final phase in the penetration testing life cycle?

In this section, you will learn to:

- Explain the types of penetration testing tools.

The key terms for this section include:

Term	Definition
White box test	Penetration test in which the ethical hacker is given full knowledge of the target or network. This test allows for a comprehensive and thorough test, but is not very realistic.

Black box test	Penetration test in which the ethical hacker has no information regarding the target or network. This type of test best simulates an outside attack and ignores the insider threats.
Gray box test	Penetration test in which the ethical hacker is given partial information of the target or network, such as IP configurations, email lists, etc. This test simulates the insider threat.
Bug bounty	These unique tests are setup by organizations such as Google, Facebook, and others. Ethical hackers can receive compensation by reporting bugs and vulnerabilities they discover.
Scope of work	A very detailed document that defines exactly what is going to be included in the penetration test. This document is also referred to as the statement of work.
Rules of engagement	A document that defines exactly how the penetration test will be carried out.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	2.2 Explain common threat vectors and attack surfaces.
	• Human vectors/social engineering
	4.3 Explain various activities associated with vulnerability management.
	• Identification methods
	  ○ Open-source intelligence (OSINT)
	• Penetration testing
	  ○ Bug bounty program
	4.4 Explain security alerting and monitoring concepts and tools.
	• Tools
	  ○ Vulnerability scanners
	5.3 Explain the processes associated with third-party risk assessment and management.
	• Agreement types
	  ○ Work order (WO)/statement of work (SOW)
	• Rules of engagement
	5.5 Explain types and purposes of audits and assessments.
	• Penetration testing
	  ○ Physical

	○ Offensive ○ Defensive ○ Integrated ○ Known environment ○ Partially known environment ○ Unknown environment ○ Reconnaissance ▪ Passive ▪ Active
TestOut Security Pro	5.2 Assessment techniques 5.2.3 Scan for vulnerabilities

7.4.1 Penetration Testing (Lesson Video)

Transcript:

Penetration testing, often referred to as "pen testing," is a practice conducted to assess the security of an IT infrastructure by safely attempting to exploit vulnerabilities. These vulnerabilities may exist in operating systems, services and application flaws, improper configurations, or risky end-user behavior. The goal of this simulated attack on a system is to identify any weak spots in an organization's defense that attackers could potentially exploit. It's designed to provide organizations with valuable insights regarding their security posture and ability to withstand cyber attacks.

Once it has been decided what exactly can be tested, a timeframe and payment agreements—if applicable— should also be established and outlined in the Scope of Work. While the Scope of Work defines what work will be done, the Rules of Engagement define exactly how that work will be carried out. It should specifically state how sensitive data should be handled and who should be notified if something unexpected happens during the test. It should also specify what test methods should be used.
In penetration testing, we distinguish between two primary testing methods: active and passive. Active Penetration Testing involves direct interaction with the system to uncover vulnerabilities. This could be attempting to exploit a known software vulnerability or trying to crack a weak password. In this method, the tester is directly engaging with the target system, often leaving a trace or log of their activities. The main goal of active penetration testing is to breach the system's defenses and evaluate the impact of such a breach.
Passive Penetration Testing, on the other hand, involves gathering information about the target system without directly interacting with it. This could involve network monitoring, analyzing system logs, or even something as simple as googling for information about the system or the organization. The goal here is to gather as much information as possible to understand the system and identify potential vulnerabilities.
Passive testing is more covert, making it less likely to be noticed and, therefore, less likely to raise any alarms.
Both methods are usually used and are important for a comprehensive penetration test because both active and passive testing methods offer unique perspectives and insights into an organization's security posture.
A pen test might involve the following steps: First, verify that a threat exists. To do so, you'd use surveillance, social engineering, network scanners, and vulnerability assessment tools to identify weak spots where vulnerabilities could be exploited.
Next is to bypass security controls. This involves looking for easy ways to attack the system. Sometimes, the simpler solutions are the most vulnerable. For example, could the network firewall be bypassed by gaining physical access to the computer in the building? Once the computer has been accessed, malware could be introduced using a USB drive.
The next step is to actively probe controls for configuration weaknesses and errors. This could include weak passwords or software vulnerabilities. Finally, once vulnerabilities have been discovered, a pen tester will prove that a vulnerability is a high risk by exploiting it to gain access to data or to install backdoors.

That's it for this lesson. In this lesson, we talked about penetration testing. We discussed outlining and documenting the goals of the pen tests before testing begins, including specifications regarding whether tests will be active or passive. We then discussed the steps that are taken during a penetration test.

7.4.2 Penetration Testing Facts

Penetration testing, also commonly referred to as pentesting or ethical hacking, is the authorized simulation of an attack against an organization's security infrastructure. This can include physical and network security.

This lesson covers the following topics:

- Types of penetration tests
- Security teams
- Documentation/contracts
- Penetration testing life cycle

Types of Penetration Tests

The purpose of a penetration test is to discover any vulnerability in an organization's network or physical security. Different types of penetration tests can be performed to simulate internal or external threats. The following table details the types of penetration tests:

Penetration Test Type	Description
Known environment (previously known as white box) testing	The penetration tester is given full knowledge of the target or network. This test allows for a comprehensive and thorough test but is unrealistic.
Unknown environment (previously known as black box) testing	The penetration tester has no information regarding the target or network. This type of test best simulates an outside attack and ignores insider threats.
Partially known environment (previously known as gray box) testing	The penetration tester is given partial information about the target or network, such as IP configurations, email lists, etc. This test simulates the insider threat.
Bug bounties	These unique tests are programs that are set up by organizations such as Google, Facebook, and many others. The organization sets strict guidelines and boundaries for ethical hackers to operate within. Any discovered vulnerabilities are reported, and the ethical hacker is paid based on the severity of the vulnerability.

Security Teams

Depending on their role, members of security operations can be placed on different teams. These teams all work together to discover and fix security vulnerabilities.
The following table describes the more common security teams:

Security Team	Description
Red team	The red team members are the ethical hackers. This team is responsible for performing the penetration tests.
Blue team	Blue team members are the defense of the system. This team is responsible for stopping the red team's advances.
Purple team	Members of the purple team work on both offense and defense. This team is a combination of the red and blue teams.
White team	The white team members are the referees of cybersecurity. This team is responsible for managing the engagement between the red and blue teams. This group typically consists of the managers or team leads.

Documentation/Contracts

Before any penetration test can take place, the goals and guidelines of the test must be established. These are spelled out in the scope of work and rules of engagement documents.
The following table describes these important documents:

Document Type	Description
Scope of work	The scope of work is a very detailed document that defines exactly what is going to be included in the penetration test. This document is also referred to as the statement of work. This document should answer the: • Who - specific IP ranges, servers, applications, etc., should be explicitly listed. • What - anything that is off limits, such as specific servers or tactics, should be explicitly listed. • When - the time frame for the penetration test. This should identify how long the test will run, the deliverables, and when the deliverables are due. • Where - the location of the penetration tester. Sometimes, the penetration tester will be located in a different state. In this case, all parties must agree on which state laws will be followed. • Why - the purpose and goals of the test. Penetration tests are often performed for compliance purposes, and these requirements must be detailed in the document. Special considerations, such as travel, required certifications, or anything else unexpected, will be defined in the scope of work. Finally, the scope of work should define payment and how to handle requests for additional work. This will help to reduce scope creep.
Rules of engagement	The rules of engagement document defines exactly how the penetration test will be carried out. The following should be defined in the rules of engagement: • Type of test - whether the test will be a white box, black box, or gray box test.

Document Type	Description
	• Data handling - an explicit statement of how sensitive data is to be handled. Be aware that the pentester will typically come across sensitive data during a penetration test. • Notifications - the detailed process on when and how to notify the IT team.

Penetration Testing Life Cycle

Once the paperwork is complete, the pentester can begin work. The following table covers the phases of the penetration testing life cycle:

Penetration Testing Life Cycle Phase	Description
Perform reconnaissance	The first phase in the pentesting process is reconnaissance, also known as footprinting. In this phase, the pentester begins gathering information on the target. This can include gathering publicly available information, using social engineering techniques, or even dumpster diving.
Scan/enumerate	Running scans on the target is the second phase. During this phase, the ethical hacker is actively engaged with the target. Enumeration is part of the scanning phase. Enumeration uses scanning techniques to extract information, such as: • Usernames • Computer names • Network resources • Share names • Running services
Gain access	The third phase takes all of the information gathered in the reconnaissance and scanning phases to exploit any discovered vulnerabilities in order to gain access. After gaining access, the pentester can perform lateral moves, pivoting to other machines on the network. The pentester will begin trying to escalate privileges with the goal of gaining administrator access.
Maintain access	Once the pentester has gained access, maintaining that access becomes the next priority. This can be done by installing backdoors, rootkits, or Trojans.
Report	The final phase is generating the test results and supporting documentation. After any penetration test, a detailed report must be compiled. Documentation provides extremely important protection for both the penetration tester and the organization.

7.4.3 Penetration Testing Methods

This lesson covers the following topics:

- Penetration tests
- Active and passive reconnaissance
- Known, partially known, and unknown testing methods
- Exercise types

Penetration Tests

A penetration test—often shortened to pen test—uses authorized hacking techniques to discover exploitable weaknesses in the target's security systems. Pen testing is also referred to as ethical hacking. A pen test might involve the following steps:

- Verify a Threat Exists — use surveillance, social engineering, network scanners, and vulnerability assessment tools to identify a vector by which vulnerabilities could be exploited.
- Bypass Security Controls — look for easy ways to attack the system. For example, if the network is strongly protected by a firewall, is it possible to gain physical access to a computer in the building and run malware from a USB stick?
- Actively Test Security Controls — probe controls for configuration weaknesses and errors, such as weak passwords or software vulnerabilities.
- Exploit Vulnerabilities — prove a vulnerability is a high risk by exploiting it to access data or install backdoors.

The critical difference from passive vulnerability assessment is that an attempt is made to actively test security controls and exploit any vulnerabilities discovered. Pen testing is an intrusive assessment technique. For example, a vulnerability scan may reveal that an SQL Server has not been patched to safeguard against a known exploit. A penetration test would attempt to use the exploit to perform code injection and compromise the server. This provides active testing of security controls.

Active and Passive Reconnaissance

Active and passive reconnaissance provides crucial information that helps penetration testers understand target systems and identify potential vulnerabilities to plan an attack effectively. A combination of active and passive reconnaissance techniques yields the most comprehensive information regarding the target environment during a penetration testing engagement.

Active reconnaissance involves probing and interacting with target systems and networks to gather information. Active reconnaissance includes activities that generate network traffic by directly requesting information from target systems. Active reconnaissance aims to discover and obtain information about the target infrastructure, services, and potential vulnerabilities. Standard techniques used in active reconnaissance include the following:

- Port Scanning — scanning a target network to identify open ports and their services.
- Service Enumeration — interacting with identified services to gather information about their versions, configurations, and potential vulnerabilities.
- OS Fingerprinting — attempting to identify the operating system running on target machines by analyzing network responses and behavior.
- DNS Enumeration — gathering information about the target's DNS infrastructure, such as domain names, subdomains, and IP addresses.
- Web Application Crawling — exploring web applications to identify pages, directories, and potential vulnerabilities.
- Open-Source Intelligence (OSINT) Gathering — collecting publicly available information from various sources like search engines, social media, public databases, and websites.
- Network Traffic Analysis — monitoring network traffic to identify patterns, devices, IP addresses, and potential vulnerabilities without actively generating traffic.
- Social Engineering — gathering information through social engineering techniques, such as deceiving employees and vendors to extract sensitive information or access credentials.

Passive reconnaissance helps penetration testers gather initial information on a target's digital footprint. It is less intrusive and has a lower detection risk than active reconnaissance techniques.

Known, Partially Known, and Unknown Testing Methods

The decision to use a known environment, partially known environment, or unknown environment penetration test is influenced by several factors, such as knowledge regarding the target system or network, the organization's risk appetite, and compliance requirements. Budget and resource constraints may also contribute to selecting the penetration testing method, as known environment testing generally requires fewer resources than partially known or unknown environment testing. The objectives of the penetration test influence the choice, with known environment testing suitable for assessing known vulnerabilities and partially known or unknown environment testing preferred for identifying unknown vulnerabilities. The complexity of the target system or network is also a factor, as more complex systems may necessitate more comprehensive testing methods. Organizations often combine different methods to achieve other objectives.

Exercise Types

Penetration testing is a crucial component of cybersecurity assessments that involves simulating real-world attacks on computer systems, networks, or applications to identify vulnerabilities and weaknesses. Different types of penetration tests exist to address specific objectives related to a security evaluation, such as testing specific systems, assessing incident response capabilities, measuring the effectiveness of physical controls, and many other areas. Different types of penetration tests allow organizations to use a flexible and prioritized approach toward security assessment.

- Physical penetration testing, or physical security testing, describes assessments of an organization's physical security practices and controls. It involves simulating real-world attack scenarios to identify vulnerabilities and weaknesses in physical security systems, such as access controls, surveillance, and perimeter defenses. Physical penetration testing aims to assess the effectiveness of physical security controls and identify potential entry points or weaknesses that an attacker could exploit. During physical penetration testing, a skilled tester attempts to gain unauthorized physical access to restricted areas, sensitive information, or critical assets within the organization using techniques like social engineering, tailgating, lock picking, bypassing alarms or surveillance systems, and exploiting physical vulnerabilities.
- Integrated penetration testing refers to a holistic approach that combines different types of penetration testing methodologies and techniques to assess the overall security of an organization's systems, networks, applications, and physical infrastructure. Integrated penetration testing aims to provide a comprehensive and realistic evaluation of an organization's security operations. The importance of integrated penetration testing lies in its ability to accurately represent the organization's security posture and identify potential risks often overlooked when testing in isolated areas. For example, offensive and defensive penetration testing comprehensively assesses an organization's security posture. Offensive testing identifies vulnerabilities and weaknesses, while defensive testing evaluates the organization's ability to detect and respond to threats. By integrating both approaches, organizations can improve their security capabilities to better protect against different threats.

Continuous pen testing is a similar concept, which focuses on technical vulnerabilities and is often configured to leverage automation, especially for CI/CD environments. Review the following for more information:
https://informer.io/resources/continuous-penetration-testing

7.4.4 Exploring Penetration Testing Tools (Demo Video)

Transcript:

In this demonstration, we're going to look at penetration testing tools that you can use to evaluate the security of your network or a particular host on your network.

As you probably know, Linux is a very popular platform for testing network and host security. First, you need to choose a Linux distribution that you can use for penetration testing. A good place to start looking for those distributions is the DistroWatch website. Many of the Linux distributions here can be run as a live CD or installed on the hard drive. A live CD is an optical disc or a bootable USB drive that has the Linux operating system installed on it. It can also have many of the security tools you need to perform a penetration test.

Because it's installed on a USB drive or an optical disc, you can insert that into the computer and boot the system off the disc. When you do, you'll have a Linux operating system up and running with the tools you need for testing.

There are several advantages associated with testing this way. First, there's a wide variety of free penetration testing tools available for the Linux operating system, and if you're using a LiveCD, you don't have to install an operating system.

If you're booting off an optical disc, there's no way for malware or anything else to actually affect the files on the disc.

There are many different distributions available. You can see their names here. You can also see the purpose of the distribution listed over here. Since we're interested in security, let's go over here and look at this one.

The most popular and well-accepted distribution for security and penetration testing is Kali Linux. Let's go ahead and click on that and see what we can learn. It tells us where the home page is and a lot of other information. I have its home page open in another tab, up here. We can read a little bit more and see what tools are actually included on this distribution. I'm going to go to the Download page. I see the latest version right here, Kali Linux 64-Bit. Over here, I can see the checksum, or hash value.

Now, I already have a copy of Kali downloaded and ready to go. We'll get to that in a minute. But I want to go back to DistroWatch. I'll go back to the previous page. I just want to point out that there's another Linux distribution called Parrot Linux. Parrot is a distribution with a collection of various utilities that are popular with penetration testers and computer forensic professionals. Here's the link to the home page, but I already have a tab opened up. Down here, you can read more about the project and learn about the different tools. Now back to Linux.

My Kali Linux ISO image is downloaded, and I've booted it up. Now it's asking me now if I want to install it to disc or just use it as a LiveCD. I'll go ahead and say, "Sure, go ahead and just use it as a LiveCD."

We'll let it launch, and pretty soon, we'll see the graph called User Interface. Each Linux distribution has its own set of package tools. This one specializes in security tools.

Here are some of the tools that Kali Linux has packaged with its distribution. Over here, you'll see the Metasploit Framework Armitage (which is a graphical user interface for that framework), Burp Suite, BeEF Cross-Site Scripting Framework, and some others that we could use to do our penetration testing.

Now, obviously, we don't have time in this short video to really discuss penetration testing or even look at a specific tool in-depth, but I'll show you how you could penetration test using one of the pre-packaged tools that we can launch from a LiveCD.

Let's start the Metasploit Framework by clicking on it. That's going to create the database where we're going to save some of the information about hosts and other things. Once that's done, I'll go ahead and bring up Armitage, which is the graphical user interface that we can use to easily interact with the extensive commands that exist as part of the Metasploit framework.

It looks like that's launched now, and we'll go ahead and launch Armitage. There it is. We'll connect, and it asks us if we want to start the remote procedure called Server. We'll say, "Sure, that'll be great. It'll connect us up to the database." It takes just a minute to make those connections.

It looks like it made the connection, and it'll launch here shortly. We're greeted with several different windows and several different options. Essentially, what we need to do is tell Armitage (and, in conjunction, the Metasploit Framework) which hosts we want to actually try to launch an attack on. We could use an Nmap scan to import that information. We could also use the Metasploit Framework scan, which will go out and ping the different machines in a certain subnet and gather information.

In this case, I'm going to keep it simple so we don't scan all the hosts on our network, and I'll just add a single host, a vulnerable Linux virtual machine that I have on the network. I know it's on the IP ending in .102. Now click Add and then OK. It says that it added it, and we've identified a host. However, we don't have much information about it yet.

But if I right-click on that host, I can scan it. Now it'll look for all the open ports, and you can see a variety of open ports that are coming up. We can also request additional information about services that might be running and on specific ports.

So you can see here that we have certain ports, but we don't know much about the services quite yet.

If we wanted to, we could try to gather some additional information. Let's do an MSF scan on that host, and that will try to identify some information about the operating system that's running on it. You'll notice that every time we run a command it launches a separate window down here, so you could go along and close these out as you go.
We know some information about this computer--some of the ports that are open, some of the services that may be running, etc. It's not a lot of data, but Metasploit keeps track of the vulnerabilities based on the ports that are open, the services that are running, the operating system that's running, and so on. You can come up here, and you can actually find attacks specific to the host that you've already discovered. We'll let it go through its database and find specific exploits that it might want to try.
So, that's complete. Now we could go through the attack menu and look for a specific exploit. You'll see here that, based on the ports that it found, it said, "Hey, FTP's open. Go ahead and try these exploits."
Each of these exploits will run the specific exploit when you click on them, and you have to provide specific values for the exploit. We don't actually know if any of these attacks will run, so the ability to check on that is built into the Metasploit framework. Some of the exploits have the ability to use this check; others are older and don't.
Let's try doing a check on some of these right now. It says this one doesn't support the check, this specific vulnerability is not exploitable, not exploitable, doesn't support the check, and so on. For the ones that don't support the check, we'd have to go through and try each of them individually.
Now, generally, as you're looking at an organization and trying to compromise hosts within it, you don't want to make a lot of noise or cause a lot of traffic on the network, so you try these very specific attacks based on what you discover. However, there is an option inside of Armitage (and, subsequently, Metasploit) that lets you do what's known as the Hail Mary, and that is just try out every single attachment possible. In the interest of time, rather than going through every single one of these attacks in a systematic way, we'll do the noisy attack because I want to show you how you can compromise the mission, and we'll see which vulnerabilities actually exist.
If we click on Hail Mary here, then it says, "Hey, are you sure you want to do this? There's going to be a flood of exploits, and it's kind of noisy. We'll say yes, that's what we intend to do, and then it'll go through that database of exploits and try to run them with a variety of different payloads so that it's trying to establish a connection to that remote machine. We'll let it do its thing here for a minute.
It went through the database, and now it's launching each of those exploits. You'll see that the exploits are going to certain ports based on the services that are running on those. We'll give it another minute here, and it'll start taking advantage of some of those exploits.
Oh look, our icon changed. That means we actually have a connection. And once it made the connection, it gathered some additional information and says, "Yep, this is definitely a Linux box." We can see what version of Linux and some other information, too.
Once we have a session, we can gather all sorts of data. We can get the dump of the passwords and the hashes associated with the passwords and then try to crack those. Maybe we compromise the accounts and get more straightforward access, but there's all of these different sessions that we could establish based on the different exploits that we just ran. Again, usually, you wouldn't do the Hail Mary because it's too noisy. But in this case, it's a quick and easy demonstration of the vulnerabilities that exist on this specific machine.
If we scroll down here, we'll see a variety of exploits that exist for this machine, and we'll give ourselves a little bit more space. You can see, as we scroll up, that it looks like there's a PHP exploit, PHP CGI injection. There's a user map, a Samba exploit, another Samba exploit, and several others that actually allowed us to have a session over to that machine.
Let's go ahead and open a new console session. I'll adjust these windows so we can see better. We can run the sessions command, and you'll see that we currently have four sessions connecting from our machine. I know that's this IP address to this vulnerable machine, right there, through PHP, and then through some command lines.
Let me go ahead and connect to one of these sessions interactively. We'll choose Session 2, and now I've got a Linux command prompt. I can do things like list the file system. I can say, "Who am I?" If you look at the bottom here, I have root access. With that access, I could grab the hashes for the passwords. I could execute any commands. I could set up additional back doors. I could do all sorts of things.
Penetration testing is awesome, and ultimately, we're trying to compromise systems so that we know how to lock them down. To do that, we often use these pre-packaged tools that come with the distribution. Kali Linux along with Metasploit and Armitage already pre-packaged is a fantastic tool to use for penetration testing.
That's it for this demonstration. In this demo, we talked about penetration testing tools, and then we talked about the concept of a LiveCD and looked at a couple of Linux distributions that you could download and use in your testing.

7.4.5 Practice Questions (Section Quiz)

q_pene_test_access_secp8

Which step in the penetration testing life cycle is accomplished using rootkits or Trojan horse programs?

Answers:

- Gain access
- Reconnaissance
- ***Maintain access**
- Enumeration

Explanation:

Once a penetration tester has gained access, maintaining that access becomes the next priority. This can be done by installing backdoors, rootkits, or Trojans.

Gain access is the third phase of the penetration test life cycle and uses the information gathered in earlier phases to exploit discovered vulnerabilities.

Reconnaissance is the first phase in the penetration testing process. This is when the penetration tester begins gathering information.

Enumeration is the second phase in the penetration testing process. The penetration tester uses scanning techniques to extract information such as usernames and computer names.

q_pene_test_blue_01_secp8

You have been hired as part of the team that manages an organization's network defense.

Which security team are you working on?

Answers:

- Purple
- Red
- ***Blue**
- White

Explanation:

Blue team members are the defense of the system. This team is responsible for stopping the red team's advances.

Members of the purple team work on both offense and defense. This team is a combination of the red and blue teams.

The red team members are the ethical hackers. This team is responsible for performing the penetration tests.

The white team members are the referees of cybersecurity. This team is responsible for managing the engagement between the red and blue teams. This group typically consists of the managers or team leads.

q_pene_test_blue_02_secp8

A recent attack on the company involving a threat actor from another country prompted the security team to host regular penetration testing exercises.

The attack involved the IT team and human resources because the breach occurred on an employee desktop.

In the upcoming training, what role would the human resource team portray along with the IT team to simulate the recent attack and its experiences?

Answers:

- ***Blue team**
- Purple team
- White team
- Red team

Explanation:

The blue team is one of two competing teams in a penetration testing exercise. The blue team performs a defensive role by operating, monitoring, and alerting controls.

The purple team members act as facilitators during a purple team exercise. This type of exercise involves collaboration between red and blue teams during breaks throughout the training.

The white team is responsible for setting the engagement rules and monitoring the penetration testing exercise.

The red team is one of two competing teams in a penetration testing exercise. The red team performs the offensive role to try to infiltrate the target.

q_pene_test_bug_bounty_secp8

As part of a special program, you have discovered a vulnerability in an organization's website and reported it to the organization. Because of the severity, you are paid a good amount of money.

Which type of penetration test are you performing?

Answers:

- White box
- Black box
- Gray box
- ***Bug bounty**

Explanation:

Bug bounties are unique tests that are set up by organizations such as Google and Facebook. The organization sets strict guidelines and boundaries for ethical hackers to operate within. Discovered vulnerabilities are reported, and the ethical hacker is paid based on the severity of the vulnerability.

In a white box test, the ethical hacker is given full knowledge of the target or network. This test is comprehensive and thorough, but it is not very realistic.

In a black box test, the ethical hacker has no information regarding the target or network. This type of test best simulates an outside attack and ignores insider threats.

In a gray box test, the ethical hacker is given partial information about the target or network, such as IP configurations and email lists. This test simulates an insider threat.

q_pene_test_recon_secp8

Which phase or step of a security assessment is a passive activity?

Answers:

- Enumeration
- ***Reconnaissance**
- Privilege escalation
- Vulnerability mapping

Explanation:

Reconnaissance is the only step of a security assessment (penetration test) that is passive.

Enumeration, vulnerability mapping, and privilege escalation are all active events in a security assessment.

q_pene_test_red_secp8

Which team performs the offensive role in a penetration exercise?

Answers:

- ***Red team**
- Blue team
- White team
- Purple team

Explanation:

The red team performs the offensive role to try to infiltrate the target. This team is one of two competing teams in a penetration testing exercise.

The blue team performs a defensive role by operating, monitoring, and alerting controls. This team is one of two competing teams in a penetration testing exercise.

The white team is responsible for setting the engagement rules and monitoring the penetration testing exercise.

The purple team members act as facilitators during a purple team exercise. This type of exercise involves collaboration between red and blue teams during breaks throughout the training.

q_pene_test_roe_01_secp8

The IT security team at a large company is reviewing its security practices to improve resilience against cyber threats. The team wants to assess the network's vulnerability to potential attacks.

The IT team decides to conduct a penetration testing exercise and hires an external cybersecurity firm with expertise in penetration testing. The goal is to identify security weaknesses and gaps that malicious actors could exploit. The company also ensures that the proper rules of engagement (ROE) are in place for the testing.

What is the primary purpose of having ROE in a penetration testing exercise?

Answers:

- ***To define the scope, limitations, and legal boundaries of the testing.**
- To determine the specific vulnerabilities present in the network.
- To identify the external cybersecurity firm responsible for conducting the test.
- To ensure the cybersecurity firm uses specialized tools for the testing.

Explanation:

ROE is essential in penetration testing exercises since it defines specific objectives, scope, limitations, and legal testing boundaries. ROE helps prevent unauthorized actions and ensures the testing remains within legal and ethical boundaries.

Penetration testing identifies specific vulnerabilities in the network. This is not the primary purpose of having ROE.

The presence of ROE does not concern identifying the external cybersecurity firm responsible for conducting the penetration test.

Specialized tools for testing are an important aspect of conducting a penetration test, but it is not the primary purpose of having ROE.

q_pene_test_roe_02_secp8

In a cybersecurity firm, the IT department is preparing for a penetration testing engagement to assess the organization's security posture. The team has decided to conduct an external penetration test on the company's public-facing web applications and networks.

The primary goal is to identify vulnerabilities and potential entry points for attackers. To ensure a smooth testing process and avoid misunderstandings, the IT team has collaborated with the company's management and relevant stakeholders to establish the assessment's rules of engagement (ROE).

What is the purpose of establishing ROE in a penetration testing engagement?

Answers:

- ***To define the scope of the assessment, testing methods, and timeframe for conducting the test.**
- To ensure the penetration test results are shared with external parties to strengthen collaboration.
- To allow penetration testers unrestricted access to all systems and data within the organization.
- To eliminate all security vulnerabilities identified during the testing process.

Explanation:

Establishing rules of engagement (ROE) in a penetration testing engagement is essential to define the assessment's scope, testing methods, and timeframe. The IT team uses ROE to prevent unintended disruptions to critical services and ensure a smooth testing process.

Although collaboration with external parties is important, the primary purpose of ROE is to establish the parameters and scope of the testing.

During the penetration test, ROE imposes limitations on accessing certain resources to prevent unauthorized access to sensitive information or systems.

The main goal of penetration testing is to identify vulnerabilities, not eliminate them. The focus is on discovering weaknesses for remediation rather than fixing them during the assessment.

q_pene_test_roe_03_secp8

The IT department in an accounting firm is gearing up for an external penetration testing engagement to evaluate the organization's security readiness.

To guarantee a seamless testing process and prevent misunderstandings, the IT team has worked closely with the company's management and relevant stakeholders to set up the rules of engagement (ROE) for the assessment.

What is the purpose of establishing rules of engagement during a penetration testing engagement?

Answers:

- ***To define the scope of the assessment, testing methods, and timeframe for conducting the test.**
- To ensure the penetration test results are shared with external parties to strengthen collaboration.
- To allow penetration testers unrestricted access to all systems and data within the organization.
- To eliminate all security vulnerabilities identified during the testing process.

Explanation:

Establishing ROE in a penetration testing engagement is essential to define the assessment's scope, testing methods, and timeframe. The IT team uses ROE to prevent unintended disruptions to critical services and ensure a smooth testing process.

Although collaboration is important, the primary purpose of ROE is to establish the parameters and scope of the testing.

During the penetration test, ROE imposes limitations on accessing certain resources to prevent unauthorized access to sensitive information or systems.

The main goal of penetration testing is to identify vulnerabilities, not eliminate them. The focus is on discovering weaknesses for remediation rather than fixing them during the assessment.

q_pene_test_social_secp8

Which of the following activities are typically associated with a penetration test?

Answers:

- Interview employees to verify that the security policy is being followed.
- Run a vulnerability scanner on network servers.
- ***Attempt social engineering.**
- Create a performance baseline.

Explanation:

Penetration testing typically uses tools and methods that are available to attackers. Penetration testing might start with attempts at social engineering or other reconnaissance activities. This may be followed by more active scans of systems and actual attempts to access secure systems.

A vulnerability scanner checks a system for weaknesses. Vulnerability scanners typically require administrative access to a system and are performed internally. They are not done to test system security. Typically, penetration testers cannot run a vulnerability scanner unless they have gained authorized access to a system.

A performance baseline is created by an administrator to identify normal network and system performance. Auditing might include interviewing employees to make sure that security policies are being followed.

q_pene_test_sow_secp8

Which of the following is a very detailed document that defines exactly what will be included in the penetration test?

Answers:

- Rules of engagement
- ***Scope of work**
- Goals and guidelines
- Payment terms

Explanation:

A scope of work is a very detailed document that defines exactly what will be included in the penetration test. This document is also referred to as the statement of work.

The rules of engagement document defines exactly how a penetration test is to be carried out.

Goals and guidelines are not a document type. The scope of work and rules of engagement documents detail the goals and guidelines of a penetration test.

Payment terms are not a document type. Payment terms are defined in the scope of work document.

q_pene_test_testing_01_secp8

Which of the following uses hacking techniques to discover internal vulnerabilities proactively?

Answers:

- Reverse engineering
- ***Penetration testing**
- Inbound scanning
- Passive reconnaissance

Explanation:

Penetration testing is the practice of testing systems and policies proactively for vulnerabilities. This approach seeks to identify vulnerabilities internally before a malicious individual can take advantage of them. Common techniques are identical to those used by hackers and include network/target enumeration and port scanning.

q_pene_test_testing_02_secp8

What is the primary purpose of penetration testing?

Answers:

- ***To test the effectiveness of your security perimeter.**
- To evaluate newly deployed firewalls.
- To assess the skill level of new IT security staff.
- To infiltrate a competitor's network.

Explanation:

The primary purpose of penetration testing is to test the effectiveness of your security perimeter. Only by attempting to break into your own secured network can you be assured that your security policy, security mechanism implementations, and deployed countermeasures are effective. It is important to obtain senior management's approval before starting a penetration test or vulnerability scanning project. Often, penetration testing or vulnerability scanning is performed by an external consultant or security outsourcing agency hired by your organization.

q_pene_test_white_box_secp8

You have been hired to perform a penetration test for an organization. You are given full knowledge of the network before the test begins.

Which type of penetration test are you performing?

Answers:

- ***White box**
- Black box
- Gray box
- Bug bounty

Explanation:

In a white box test, the ethical hacker is given full knowledge of the target or network. This test allows for a comprehensive and thorough test, but it is not very realistic.

In a black box test, the ethical hacker has no information regarding the target or network. This type of test best simulates an outside attack and ignores insider threats.

In a gray box test, the ethical hacker is given partial information about the target or network, such as IP configurations, email, and lists. This test simulates an insider threat.

Bug bounties are unique tests that are set up by organizations such as Google and Facebook. The organization sets strict guidelines and boundaries for ethical hackers to operate within. Any discovered vulnerabilities are reported, and the ethical hacker is paid based on the severity of the vulnerability.

q_pene_test_white_secp8

You have been promoted to team lead of one of the security operations teams.

Which security team are you now a part of?

Answers:

- Purple
- Red
- Blue
- ***White**

Explanation:

The white team members are the referees of cybersecurity. This team is responsible for managing the engagement between the red and blue teams. This group typically consists of the managers or team leads.

Blue team members are the defense of the system. This team is responsible for stopping the red team's advances.

Members of the purple team work on both offense and defense. This team is a combination of the red and blue teams.

The red team members are the ethical hackers. This team is responsible for performing the penetration tests.

q_pentest_meth_active_reconnaissance_secp8

You are a cybersecurity specialist conducting an active reconnaissance as part of a penetration test for a client. Your goal is to gather as much information as possible about the client's network without raising any alarms.

Which of the following techniques would be MOST effective in achieving this goal?

Answers:

- Social engineering
- ***DNS enumeration**
- Network traffic analysis
- Physical intrusion

Explanation:

DNS enumeration is the correct answer. DNS enumeration involves gathering information about the client's DNS infrastructure, such as domain names, subdomains, and IP addresses. This technique is a part of active reconnaissance and can provide valuable information about the client's network without raising alarms.

While social engineering can be a powerful tool in a penetration tester's arsenal, it is not the best choice in this scenario. Social engineering involves manipulating individuals to reveal confidential information, which can be risky and potentially alert the client to your activities.

Network traffic analysis is typically considered a passive reconnaissance technique, not an active one. It involves monitoring network traffic to identify patterns, devices, IP addresses, and potential vulnerabilities without actively generating traffic.

Physical intrusion involves gaining unauthorized physical access to a location, which is not only illegal but also likely to raise alarms. It is not a part of active reconnaissance and is not the best choice in this scenario.

q_pentest_meth_defensive_secp8

A financial institution has hired a team of security professionals to evaluate its defensive security measures and incident response procedures.

The institution wants to assess how effectively it can detect and respond to cyber threats and identify areas for improvement. However, it is not looking for an assessment of the digital vulnerabilities or physical security at this time.

What type of penetration testing is the MOST suitable for this financial institution?

Answers:

- Offensive penetration testing
- ***Defensive penetration testing**
- Physical penetration testing
- Integrated penetration testing

Explanation:

Defensive penetration testing, or "blue teaming," assesses an organization's defensive security measures, detection capabilities, and incident response procedures. This method is the most suitable type of penetration testing for the financial institution, as it is specifically interested in evaluating its ability to detect and respond to cyber threats.

Offensive penetration testing, or "red teaming," involves simulating real-world cyberattacks on an organization's digital systems to identify vulnerabilities and weaknesses.

Physical penetration testing evaluates an organization's physical security practices and controls to identify potential entry points or weaknesses that an attacker could exploit.

Integrated penetration testing combines different types of penetration testing methodologies and techniques to evaluate an organization's security operations comprehensively.

q_pentest_meth_integrated_secp8

A multinational corporation hires a team of penetration testers to assess the security of its operations across both digital and physical domains.

The team incorporates a few penetration testing techniques, including simulating real-world cyberattacks, evaluating defensive measures, and assessing physical security controls.

What approach to penetration testing is the corporation using?

Answers:

- Offensive penetration testing
- Defensive penetration testing
- Physical penetration testing
- ***Integrated penetration testing**

Explanation:

Integrated penetration testing is a holistic approach that combines different types of penetration testing methodologies and techniques to assess the overall security of an organization's systems, networks, applications, and physical infrastructure.

Offensive penetration testing, or "red teaming," is a proactive approach to simulate real-world cyberattacks on an organization's systems, networks, and applications to identify vulnerabilities and weaknesses.

Defensive penetration testing, or "blue teaming," evaluates an organization's defensive security measures, detection capabilities, and overall resilience against cyber threats.

Physical penetration testing assesses an organization's physical security practices and controls. It involves simulating real-world attack scenarios to identify vulnerabilities and weaknesses in physical security systems.

q_pentest_meth_known_secp8

A cybersecurity team is preparing to conduct a comprehensive security assessment. The team has access to system documentation, network diagrams, and source code and has permission to interview IT staff.

What type of testing environment is the team operating within?

Answers:

- ***Known environment**
- Partially known environment
- Unknown environment
- Uncontrolled environment

Explanation:

A known environment refers to situations where the testers have complete information about the system's internals, design, infrastructure, and security measures.

A partially known environment refers to situations where the tester only has limited information about the system. However, in this scenario, the team has full access to detailed system information, going beyond what would typically be available in a partially known environment.

The tester has little to no information about the system in an unknown environment. Given the extensive information available to the team in this scenario, this environment is still unknown.

An uncontrolled environment is not a standard term used in system testing.

q_pentest_meth_offensive_secp8

A cybersecurity team at an organization prepares to carry out an assessment that aims to mimic potential attackers' tactics, techniques, and procedures (TTPs) to identify vulnerabilities and weaknesses in the organization's digital systems.

What type of penetration test is the team about to conduct?

Answers:

- ***Offensive penetration testing**
- Defensive penetration testing
- Physical penetration testing
- Integrated penetration testing

Explanation:

Offensive penetration testing is a proactive and controlled approach to simulate real-world cyberattacks on an organization's systems, networks, and applications to identify vulnerabilities, weaknesses, and potential attack vectors that malicious actors could exploit.

Defensive penetration testing evaluates an organization's overall resilience against cyber threats, not actively trying to find vulnerabilities as an attacker might.

Physical penetration testing involves assessing an organization's physical security practices and controls, such as access controls, surveillance, and perimeter defenses.

While integrated penetration testing can include offensive penetration testing, the scenario specifically describes an offensive penetration test.

q_pentest_meth_partially_01_secp8

The IT security team of a company has concerns about network vulnerabilities and hires an external penetration tester to evaluate its security controls and identify potential risks.

The company provides the penetration tester with fragments of network information and permits them to use reconnaissance techniques for further information gathering.

What penetration testing method is the company using?

Answers:

- Known environment penetration testing
- ***Partially known environment penetration testing**
- Unknown environment penetration testing
- Open-source intelligence gathering

Explanation:

Partially known environment penetration testing involves the tester's limited knowledge about the target system or network. During the test, the tester may employ reconnaissance techniques to gather additional information about the target.

Known environment penetration testing involves the tester having detailed knowledge about the target system or network, including information about the network architecture, hardware and software configurations, system vulnerabilities, and users.

Unknown environment penetration testing involves the tester having little knowledge about the target system or network. This testing mimics a scenario where an attacker has no preexisting information about the target infrastructure.

Open-source intelligence gathering is a passive reconnaissance technique that collects publicly available information from various sources. It is not a penetration testing method.

q_pentest_meth_partially_02_secp8

A cybersecurity team is investigating a complex cyber threat landscape for a large financial institution. The team is aware of some potential threats due to previous encounters and security measures in place, but the evolving nature of the landscape presents new threats and challenges.

What type of cyber environment is the team dealing with?

Answers:

- ***Partially known environment**
- Reconnaissance
- Unknown environment
- Fully known environment

Explanation:

The cybersecurity team is aware of some threats but not all due to the evolving nature of the cyber landscape. This level of awareness suggests they're operating in a partially known environment.

Performing reconnaissance is a technique used by cybersecurity teams to learn more about their environment and potential threats. It doesn't describe the type of environment the team is operating in.

A fully unknown environment suggests that the cybersecurity team needs to gain prior knowledge of potential threats or challenges, which isn't true in this scenario.

A fully known environment suggests that the cybersecurity team has complete knowledge of all potential threats and challenges, which needs to be more accurate in this scenario.

q_pentest_meth_passive_reconnaissance_secp8

You are a cybersecurity analyst tasked with gathering information about a potential target for a penetration test. Your goal is to collect as much data as possible without alerting the target or leaving any trace of your activities.

Which reconnaissance technique would be MOST suitable for this task?

Answers:

- Active reconnaissance
- ***Passive reconnaissance**
- Social engineering
- Physical intrusion

Explanation:

Passive reconnaissance is the correct answer. Passive reconnaissance involves collecting information about a target without directly interacting with the target system or network. This can include activities like observing publicly available information, analyzing metadata, or monitoring network traffic. Passive reconnaissance is less likely to alert the target or leave traces of your activities, making it the most suitable technique for this task.

While active reconnaissance can provide valuable information about a target, it involves directly interacting with the target system or network. This can potentially alert the target to your activities or leave traces of your actions, making it less suitable for this task.

Social engineering involves manipulating individuals to reveal confidential information. While it can be effective, it also carries a risk of detection, especially if the individuals become suspicious. Therefore, it is not the most suitable technique for this task.

Physical intrusion involves gaining unauthorized physical access to a location, which is not only illegal but also likely to raise alarms. It is not a part of reconnaissance and is not suitable for this task.

q_pentest_meth_pen_test_steps_secp8

You are a cybersecurity analyst working for a large corporation. You have been tasked with conducting a penetration test to identify potential vulnerabilities in the company's network.

You have already completed the initial step of verifying that a threat exists.

What should be your next step in the penetration testing process?

Answers:

- Exploit vulnerabilities.
- ***Bypass security controls.**
- Actively test security controls.
- Report findings to management.

Explanation:

Bypass security controls is the correct answer. After verifying that a threat exists, the next step in the penetration testing process is to bypass security controls. This involves looking for easy ways to attack the system. For example, if the network is strongly protected by a firewall, is it possible to gain physical access to a computer in the building and run malware from a USB stick?

Exploit vulnerabilities is not the correct next step in the penetration testing process. Before exploiting vulnerabilities, you need to identify potential weaknesses in the security controls. This is done by bypassing the security controls and actively testing them.

While actively testing security controls is an important step in the penetration testing process, it is not the next step after verifying that a threat exists. Before you can actively test security controls, you need to bypass them to identify potential weaknesses.

Reporting findings to management is the final step in the penetration testing process, not the next step after verifying that a threat exists. Before you can report your findings, you need to bypass security controls, actively test them, and exploit any vulnerabilities you find.

q_pentest_meth_process_steps_secp8

You are a cybersecurity analyst at a large corporation and have been tasked with conducting a penetration test to identify potential vulnerabilities in the company's network.

You have just completed the reconnaissance phase and have identified a potential vulnerability in the company's SQL Server. What is your next step?

Answers:

- Immediately report the vulnerability to the company's management.
- Exploit the vulnerability to gain access to the server.
- ***Actively test the security controls around the vulnerability.**
- Bypass the security controls and gain physical access to a computer in the building.

Explanation:

Actively test the security controls around the vulnerability is the correct answer. After identifying a potential vulnerability, the next step in a penetration test is to actively test the security controls around it. This can involve probing for configuration weaknesses and errors, such as weak passwords or software vulnerabilities.

Immediately report the vulnerability to the company's management is not the correct next step in a penetration test. While it's important to keep management informed, the purpose of a penetration test is to actively test and exploit vulnerabilities to understand their potential impact. Reporting should be done after these steps have been completed.

Exploit the vulnerability to gain access to the server is not the correct next step. Before attempting to exploit a vulnerability, it's important to actively test the security controls around it. This can help identify any additional weaknesses or protections that may affect the exploitation process.

Bypass the security controls and gain physical access to a computer in the building is not the correct next step. While physical access can sometimes be a part of a penetration test, it's not the appropriate response in this scenario. The identified vulnerability is in the SQL Server, so the focus should be on testing and potentially exploiting this vulnerability.

q_pentest_meth_tester_secp8

A company hires a team of penetration testers to evaluate the security posture of its newly developed web application. After a comprehensive analysis, the testers submit their findings, detailing potential vulnerabilities.

The company's security officer reviews the report and contemplates the essential differences between how threat actors and penetration testers would exploit the identified vulnerabilities.

What distinct motive differentiates a professional penetration tester from a threat actor when it comes to exploiting vulnerabilities in a system?

Answers:

- Penetration testers aim to damage or disrupt the system.
- Threat actors provide a detailed report of their findings.
- ***Penetration testers work to identify vulnerabilities and improve system security.**
- Threat actors have permission to test the system.

Explanation:

Penetration testers are professionals who evaluate the security posture of systems and networks. Their main goal is to identify vulnerabilities and recommend solutions to bolster security.

Unlike malicious attackers (threat actors), penetration testers do not have a motive to cause harm or service disruption. Their primary objective is to discover vulnerabilities without causing damage.

While penetration testers provide a comprehensive report to the organization after their analysis, threat actors typically do not submit any findings, as their primary goal might be theft, disruption, or other malicious intent.

Threat actors do not have permission, and their actions are illegal. Conversely, penetration testers operate under explicit consent from the organization.

q_pentest_meth_testing_secp8

An organization considers a new third-party vendor to provide critical technology solutions. It is nearing the final stages of the vendor selection process and wants to ensure a robust assessment of the vendor's security practices and risk management capabilities.

Which method would be MOST suitable for the organization to gain an in-depth understanding of the vendor's internal security controls, identify potential vulnerabilities in its systems, and validate the effectiveness of its security measures?

Answers:

- Rely on the vendor's self-assessment report.
- Perform a supply chain analysis.
- ***Conduct a penetration test.**
- Request evidence of internal audits.

Explanation:

Penetration testing is a proactive and in-depth method of testing a vendor's defenses. It helps organizations discover potential vulnerabilities in the vendor's systems, networks, and applications that attackers could exploit.

Self-assessments can be biased and may not accurately disclose all potential risks or vulnerabilities.

While supply chain analysis is important, it mostly helps the company understand the risk associated with multiple entities within the organization's supply chain.

While internal audits are important for providing insight into a vendor's risk management and compliance, they might not provide a detailed understanding of specific security vulnerabilities in the vendor's systems the way a penetration test would.

q_pentest_meth_unknown_01_secp8

A software company has completed in-house testing and auditing and is bringing in an outside source to attempt to compromise the new software. The project head wants to ensure that the MOST realistic testing goes against the software.

What type of penetration testing will the outside source use on this new software?

Answers:

- ***Unknown environment**
- Known environment
- Partially known environment
- Environmental variables

Explanation:

In an unknown environment, the contractor has no information about the software and performs extensive reconnaissance to attempt an attack on the software as an external threat would.

Known environment, the contractor has all the information about the software and uses this to compromise the application. This process is an example of an insider threat to the software.

In a partially known environment test, the contractor has some information about the software but will still need to perform reconnaissance.

Environmental variables related to vulnerability analysis refer to the systems and software a company has currently in use. Age and diversity, as well as complexity, are all variables. This process is not a software test.

q_pentest_meth_unknown_02_secp8

A cybersecurity team is preparing to conduct a comprehensive security assessment. The team needs more detailed information about the system, including system documentation, network diagrams, source code, and permission to interview IT staff.

What type of testing environment is the team operating within?

Answers:

- ***Unknown environment**
- Known environment
- Partially known environment
- Uncontrolled environment

Explanation:

The tester needs more information about the system in an unknown environment. Given the lack of information available to the team in this scenario, this environment is applicable.

A known environment refers to situations where the testers have complete information about the system's internals, design, infrastructure, and security measures.

A partially known environment refers to situations where the tester only has limited information about the system. Since the team cannot access detailed system information, this could not be the environment.

An uncontrolled environment is not a standard term used in system testing or this scenario.

8.0 Network and Endpoint Security

8.1 Operating System Hardening

As you study this section, answer the following questions:

What is hardening? How does it benefit security?

How do you reduce the attack surface of a device?

Why should you install only software that you need?

What is a security baseline?

What is the difference between a hotfix and a patch? Why would you use one instead of the other?

In this section, you will learn to:

Harden an operating system.

Manage automatic updates.

Configure automatic updates.

Configure Microsoft Defender Firewall.

The key terms for this section include:

Term	Definition
Patches	A small unit of supplemental code meant to address either a security problem or a functionality flaw in a software package or operating system.
Patch management	Identifying, testing, and deploying OS and application updates. Patches are often classified as critical, security-critical, recommended, and optional.
Allow list	A security configuration where access is denied to any entity (software process, IP/domain, and so on) unless the entity appears on an allow list, also known as a whitelist.
Block list	A security configuration where access is generally permitted to a software process, IP/domain, or other subject unless it is listed as explicitly prohibited.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions Encryption Level Full-disk 2.2 Explain common threat vectors and attack surfaces Removable device 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Access control Application allow list Patching Encryption Monitoring Least privilege Configuration enforcement Decommissioning Hardening techniques Installation of endpoint protection Host-based firewall Disabling ports/protocols Default password changes Removal of unnecessary software 4.1 Given a scenario, apply common security techniques to computing resources. Secure baselines

	Establish Deploy Maintain Hardening targets Workstations 4.3 Explain various activities associated with vulnerability management. Vulnerability response and remediation Patching 4.5 Given a scenario, modify enterprise capabilities to enhance security. Operating system security Group Policy 4.6 Given a scenario, implement and maintain identity and access management. Provisioning/de-provisioning user accounts Access controls Least privilege Multifactor authentication 4.7 Explain the importance of automation and orchestration related to secure operations. Use cases of automation and scripting Security groups
TestOut Security Pro	1.1 Manage identity 1.1.1 Manage Windows local and domain users and groups 1.2 Harden authentication 1.2.5 Configure and link Group Policy Objects (GPO) 2.1 Harden physical access 2.1.3 Install and configure a firewall 3.1 Harden computer systems

	3.1.2 Configure anti-virus protection 3.1.4 Configure Windows Update 3.2 Implement application defenses 3.2.1 Implement an application allow list

8.1.1 Operating System Hardening (Lesson Video)

Transcript:

The main purpose of hardening an operating system is to eliminate as many security risks as possible.

If they're not implemented correctly, passwords can be an operating system vulnerability. To mitigate this issue, passwords should be complex. This means that they should be at least 12 characters long and include capital letters, numbers, and symbols.
In a high-security environment, passwords alone are just too vulnerable; multi-factor authentication is a must. Multi-Factor authentication uses two or more authentication factors, such as something you know, like a password or a pin; something you have, like a smart card; or something you are, like a retina scan.
Another hardening technique is to limit operating system exposure. To do this, you should first ensure that administrative privileges are only provided to system administrators. Next, system management should ensure that users are only given the rights and permissions required to do their jobs â€“ nothing more. This concept is called the principle of least privilege.
You should also ensure that only the software used by administrators and users is installed on your systems. Removing unnecessary software reduces the attack surface, giving hackers fewer chances to exploit a system.
For that reason, you should disable unnecessary or non-essential services. Unused services often aren't configured securely or configured at all. They're attractive to hackers, and sustained attacks can go unnoticed.
Examples of non-essential services are TFTP, Telnet, and SNMP. These are all clear text services.
DNS, ICMP, and NTP are essential protocols and services. If your system isn't a DNS server, you'll want to disable the DNS service. It's good to note that in Windows 10, some of these services are disabled by default, such as Telnet and TFTP.
Systems should also be protected with security software, such as antivirus, anti-spyware, anti-rootkit, and personal or client firewall software. When you use this type of security software, ensure that it comes from a trusted publisher and doesn't conflict with other programs in use.
Another aspect of hardening an operating system is to develop system baselines. A baseline is a standard applied to systems. It defines standards for configuration management, security policies, software and backup policies, and several other topics. Developing standard baselines helps ensure that an organization is working at peak efficiency and security base.
If baselines are appropriately configured and applied, it becomes fairly easy to detect systems that may be vulnerable since they're operating outside established standards. There are several types of baselines to consider.
A configuration baseline establishes the standard configuration for all systems. This baseline defines the default system configuration, including loaded software, enabled services, and domain connectivity. Further configuration and changes to the configuration can be achieved via Group Policies.
A security baseline defines the secure computing components, such as the security software used to protect individual systems, that are installed and configured on your workstations. Definition updates and changes are all configured and ready to go as part of the baseline. In addition, any enterprise hardware such as firewalls, gateways, and SPAM filters (and their configurations) are specified as part of this baseline.
A software & use baseline defines acceptable system use and supported software. Software managers can control a user's ability to employ software. This can also help administrators by keeping track of use licenses and, potentially, reduce expenditures by limiting the number of purchased software use licenses. It also helps define appropriate use policies by defining how to implement Group Policies to provide allowed services while preventing unauthorized use of disallowed services.
As a whole, baselines should be reviewed and modified to reflect the changes organizations make constantly. When new software or new versions of existing software are introduced, the baseline must be changed to reflect the

update. If security policies change, such as the frequency of password changes or password complexity requirements, the baseline must be updated again. Baselines are not static. They should be thought of as a guideline or practice to ensure the efficient and safe use of computer systems. Further, they should reduce the attack surface that threat actors can exploit to gain access to your system.

The last topic for this lesson is patch management. Patch management is the process of managing patches for all of your systems. A patch management system can provide centralized control. A common system is Microsoft's patch management product Windows Software Update Service, or WSUS for short.

Patch management activities include determining which patches are needed on your systems, applying the patches, and then auditing for the successful application of those patches. Before patches are applied, a small group of deployments is usually tested, especially for service packs. If a problem exists with the patch process, it's much easier to handle a small deployment than to face a system-wide crisis. This group is generally comprised of power users who know how to fix issues. They're given additional permissions and rights so they can continue working if something goes wrong with deployment.

Larger organizations may have a lab that represents the types of systems in use by users throughout the enterprise. Tests are conducted within the lab and verified before they're released to the general user community. A best practice is to create a system restore point before you patch systems just in case something goes wrong.

Now, sometimes, a patch is overkill, so we also have hot fixes to choose from. Hot fixes are quick fixes for problems. Usually, you don't install the hot fix unless you have the specific problem it's intended for. Hot fixes sometimes address a specific customer situation and may be distributed only to that customer. Hot fixes are also commonly used to address freshly discovered security holes.

A patch could also be a temporary quick fix. Generally, a manufacturer tests patches more than a hot fix, and patches are designed for wider deployment. Patches can include hot fixes that the manufacturers have thoroughly tested for mass deployment. Patches typically contain fixes for security holes and discovered system bugs. As an example, Microsoft releases security patches once per month.

Service packs are installable packages that include several patches from the same vendor for various applications. Service packs usually include patches and hot fixes that have been tested by the manufacturer for wide deployment. Before fully deploying the service pack, you'll usually let a small designated group test it to make sure it's fully compatible with your system. But if you have a specific issue, you may want to install a patch or a hot fix immediately.

That's it for this lesson. In this video, we talked about operating system hardening. First, we discussed the importance of using secure passwords, limiting administrative privileges, installing only needed, secure software, and removing non-essential services. Next, we discussed the role baselines play in keeping our systems secure. We ended this lesson by discussing the importance of applying the latest patches and service packs to your systems.

8.1.2 Hardening Facts

This lesson covers the following topics:

- Operating system security
- Workstations
- Installing endpoint protection
- Patch management
- OS hardening best practices
- Encryption techniques

Operating System Security

Operating system security encompasses many practices that aim to protect against unauthorized access, data breaches, malware infections, and other security threats. There are many considerations and requirements when securing an

operating system because operating systems are complicated and powerful software products operating at the core of all information systems. Many security concepts apply to operating system security, including access controls, authentication mechanisms, secure configurations, application security, secure coding, patch management, endpoint protection, user awareness training, and monitoring.

Hardening describes changing an operating system or application to make it operate securely. The need for hardening must be balanced against functional requirements and usability because hardening can often negatively impact how applications work or interoperate.

Best practice baselines play a critical role in device hardening by providing a standard set of guidelines or checklists for configuring devices securely. These baselines, often developed by cybersecurity experts or organizations, offer a starting point for systems administrators to ensure that devices are configured according to industry security standards. Many of the requirements can be applied automatically via a configuration baseline template. The essential principle is of least functionality; that a system should run only the protocols and services required by legitimate users and no more. This reduces the potential attack surface.

- Interfaces provide a connection to the network. Some machines may have more than one interface. For example, there may be wired and wireless interfaces or a modem interface. Some machines may come with a management network interface card. If any of these interfaces are not required, they should be explicitly disabled rather than simply left unused.

- Services provide a library of functions for different types of applications. Some services support local features of the OS and installed applications. Other services support remote connections from clients to server applications. Unused services should be disabled.

- Application service ports allow client software to connect to applications over a network. These should be disabled or blocked at a firewall if remote access is not required. Be aware that a server might be configured with a nonstandard port. For example, an HTTP server might be configured to use 8080 rather than 80. Conversely, malware may try to send nonstandard data over an open port. An intrusion detection system should detect network data that does not conform to the expected protocol format.

- Persistent storage holds user data generated by applications, plus cached credentials. Disk encryption is essential to data security. Self-encrypting drives can be used so that all data at rest is always stored securely.

It is also important to establish a maintenance cycle for each device and keep up to date with new security threats and responses for the particular software products that you are running.

Workstations

Workstations operate at the frontline of an organization's activities and present unique concerns regarding endpoint hardening compared to other devices. Due to the varied tasks and numerous applications associated with workstation use, they generally have a large attack surface. Hardening practices to minimize this attack surface include removing unnecessary software, limiting administrative privileges, strictly managing application installations and updates, and many other changes. Furthermore, since employees operate workstations, user-focused security strategies are essential, including regular training and awareness activities to educate users about threats such as phishing and promoting secure behaviors such as strong password practices, responsible Internet use, and careful handling of sensitive data, among other practices.

Additionally, configuring workstation settings for increased security, like automatic updates, screen locks, firewalls, endpoint protection, intrusion detection and prevention, increased logging, encryption, monitoring, and many other protections, are essential. Also, the need to secure peripheral devices like USB ports is unique to workstations. It is often achieved using features of endpoint protection software and the implementation of strict device control policies. Lastly,

given the various roles and responsibilities assigned to different workstations, segmentation is crucial to restrict communications and limit the potential for malware or attackers to propagate across the network.

Installing Endpoint Protection

To ensure maximum protection and efficient management, deploying and managing endpoint protection agents on workstations, laptops, and servers in an enterprise environment requires strategic planning and adherence to established best practice configuration and management practices.

Create a deployment plan with considerations such as the deployment order (such as which devices or departments get agents first), time frames, and leveraging stages to limit potential disruptions caused by endpoint protection settings.

Standardize configurations for endpoint protection across all devices to ensure consistency in protection levels and simplify compliance management.

Automate deployments using tools like Microsoft's System Center Configuration Manager (SCCM), Group Policy, or third-party solutions to save time, improve consistency, and reduce the risk of human error.

Updates and patches to endpoint protection agent software and definitions are required to ensure the highest levels of protection against the latest known threats.

Monitor endpoint protection agents to check for alerts or signs of potential security incidents, verify that agents are running, and ensure that updates and patches are applied successfully.

Centralize management to provide a comprehensive view of endpoint configurations, updates, and status. Centralized management also allows administrators to enforce global security policies.

Patch Management

No operating system, software application, or firmware implementation is free from vulnerabilities. As soon as a vulnerability is identified, vendors will try to correct it. At the same time, attackers will try to exploit it. Automated vulnerability scanners can effectively discover missing patches for the operating system, plus a wide range of third-party software apps and devices/firmware. Scanning is only useful if effective procedures are i n place to apply the missing patches.

In residential and small networks, hosts are typically configured to auto-update, meaning they check for and install patches automatically. The major OS and applications software products are well supported in terms of vendor-supplied fixes for security issues. In Windows, this process is handled by Windows Update, while in Linux, it can be configured via yum-cron or apt unattended upgrades , depending on the package manager used by the distribution.

There can also be performance and management issues when multiple applications run update clients on the same host. For example, as well as the OS updater, there is likely also a security software update, browser updater, Java updater, OEM driver updater, and so on. These issues can be mitigated by deploying an enterprise patch management suite. Some suites, such as Microsoft™ System Center Configuration Manager (SCCM)/Endpoint Manager (docs.microsoft.com/en-us/mem/configmgr), are vendor-specific, while others are designed to support third-party applications and multiple OSes.

Testing patches before deploying them into the production environment is crucial for maintaining the stability and security of software. By conducting thorough testing, organizations can identify potential issues or conflicts arising from the patch, ensuring that it does not introduce new vulnerabilities or disrupt critical operations. Testing helps mitigate the risk of unintended consequences and facilitates a more controlled deployment process, ultimately safeguarding the integrity and

reliability of the environment. Testing is typically performed in testing environments built to mirror the production environment as much as appropriate.

Patch management can be difficult for legacy systems, proprietary systems, and systems from vendors without robust security management plans, such as some types of Internet of Things devices. These systems will need compensating controls or some other form of risk mitigation if patches are not readily available.

OS Hardening Best Practices

Different hardening approaches are required to protect endpoints in response to varied and constantly evolving cybersecurity threats. These threats require a layered and comprehensive defense strategy addressing vulnerabilities at multiple levels, from physical access to network protocols, operating system configurations, and user behaviors.

Technique	Definition
Access controls	Access control refers to regulating and managing the permissions granted to individuals, software, systems, and networks to access resources or information. Access controls ensure that only authorized entities can perform specific actions or access certain data, while unauthorized entities are denied access. Access control concepts apply to networks, physical access, data, applications, and the cloud. *Access control lists* (ACLs) in computer systems and networks are used to enforce access control policies. An ACL is a list of rules or entries that specify which users or groups are allowed or denied access to specific resources or perform certain actions. In networks, ACLs are associated with routers, firewalls, or similar devices and define rules that determine how network traffic is filtered or forwarded based on criteria like source IP addresses, destination IP addresses, ports, or protocols. ACLs can help to control network access and protect against unauthorized or malicious activities. ACLs control access to files, directories, or system resources in operating systems and file systems. Each access control entry (ACE) typically contains a user or group identifier and associated permissions controlling actions that are allowed or denied. These permissions often include read, write, execute, and sometimes more granular limits such as modify, delete, or list. While ACLs offer flexibility and control, managing complex access control policies with numerous ACL entries can become challenging. Complexity increases the risk of misconfigurations. Therefore, proper planning, periodic reviews, and best practice configurations are essential when implementing and maintaining ACLs.
Principle of least privilege	Implementing the *principle of least privilege* (PoLP) is a cornerstone of improving endpoint protection and minimizing the risk of security issues. The principle of least privilege dictates that users, applications, and processes should only be granted the minimum permissions necessary to complete their duties and nothing more. There are several practical methods for implementing least privilege. An essential first step to effectively implementing least privilege is thoroughly auditing user roles, privileges, and responsibilities. This process allows organizations to understand what access each user needs to perform their job role effectively. Access controls and permissions can be adjusted to adopt a principle of least privilege that best reflects the audit results.

	The principle of least privilege also applies to software applications and operating systems, not just to users. For instance, ensuring that applications run with the minimum necessary permissions can prevent them from being exploited to carry out privileged actions.
Application allow lists and block lists	One element of endpoint configuration is an execution control policy that defines applications that can or cannot be run. An allow list (or approved list) denies execution unless the process is explicitly authorized. A block list (or deny list) generally allows execution but explicitly prohibits listed processes. The contents of allow lists and block lists need to be updated in response to incidents and ongoing threat hunting and monitoring. Threat hunting may also provoke a strategic change. For example, if you rely principally on explicit denies, but your systems are subject to numerous intrusions, you will have to consider adopting a "least privileges" model and using a deny-unless-listed approach. This sort of change can be highly disruptive, however, so it must be preceded by a risk assessment and business impact analysis.
Monitoring	Monitoring plays a vital role in endpoint hardening, helping to enforce and maintain the security measures put in place during the hardening process. Once devices are hardened, monitoring helps to ensure these conditions remain in place. Security analysts can detect changes that weaken the hardened configuration through continuous monitoring. For instance, if a previously disabled port is detected as open or a service that was disabled is changed to enabled, monitoring tools can alert analysts of the change, which may indicate a breach. Additionally, monitoring can provide valuable data for compliance and auditing purposes. Regular reports on the status of endpoint devices can verify that hardening baselines have been effectively deployed and maintained, supporting compliance with various regulations and industry standards.
Configuration enforcement	Configuration enforcement describes methods used to ensure that systems and devices within an organization's network adhere to mandatory security configurations. Configuration enforcement generally depends upon a few important capabilities. **Standardized configuration baselines** are defined by organizations like NIST, CIS, or the organization itself and used as the benchmark for how systems and devices should be configured. **Automated configuration management tools** are used to apply and maintain standardized configuration baselines across the environment automatically. **Continuous monitoring and compliance checks** are crucial to detect deviations from mandatory configurations.

	Change management processes ensure configuration changes are properly reviewed, tested, and approved before implementation.
Decommissioning	Decommissioning processes play a vital role in supporting security within an organization. When a device is no longer needed, it often contains residual data, potentially sensitive information, and system configurations that could be exploited. A thorough and systematic decommissioning process ensures that all data is securely erased or overwritten to reduce the risk of exposure. Decommissioning also involves resetting devices to their factory settings and eliminating any residual settings. Updating inventory records during decommissioning is also important to maintain an accurate account of active assets and support compliance requirements that mandate accurate asset tracking and secure disposal.
Changing defaults and removing unnecessary software	Changing default passwords and removing unnecessary software are two fundamental practices in hardening an endpoint to strengthen its security posture. Default passwords, often set by manufacturers, are widely known and easily discoverable, making devices that use them particularly vulnerable to unauthorized access. Therefore, changing these passwords to strong, unique credentials is crucial as part of the initial setup process for any device or system. Removing unnecessary software is another critical step when hardening an endpoint. Each software application introduces potential vulnerabilities that malicious actors could exploit. The attack surface is significantly reduced by reducing the number of software applications to only those necessary for the device's intended function. This includes removing unnecessary applications and disabling unnecessary features or services within the remaining applications. This practice simplifies the maintenance and patch management process because there are fewer applications to update, reducing the chances of missing a critical security patch.

Encryption Techniques

Endpoint encryption is critical to protecting sensitive data, especially in an enterprise setting. Several different approaches are required to protect data on endpoints. Some important ones include the following:

Full Disk Encryption (FDE) encrypts the entire hard drive of a device. It ensures that all data, including the operating system and user files, are protected even while the operating system is not running. Tools like BitLocker for Windows and FileVault for macOS provide full disk encryption capabilities.

Removable Media Encryption ensures that data remains protected even when physically removed from devices such as SDCards or USB mass storage devices. Many FDE tools also include options for encrypting removable media.

Virtual Private Networks (VPNs) complement endpoint encryption by providing a secure tunnel for data transmission that protects against eavesdropping, on-path, and many other attack types.

Email Encryption protects sensitive information stored in emails using protocols like PGP (Pretty Good Privacy) or S/MIME (Secure/Multipurpose Internet Mail Extensions) .

8.1.3 Hardening an Operating System (Demo Video)

Transcript:

Let's take a look at some general strategies for hardening an operating system. Here we are on a domain controller running Windows Server. Let's go to Tools and go to Active Directory Users and Computers. There are a lot of options here for locking things down. The default user is the administrative user, and that's a very well-known account. It's built into the system, so a lot of people understand what it is and what it does.

One thing you might consider doing is renaming the account in Properties, or you could create a new account for the administrator. Be aware that the Administrator account has many groups assigned to it, such as domain admins, enterprise admins, and schema admins. By default, the Domain Admins group grants the rights to configure group policies.

Samantha Carter has been promoted to a junior administrator, so we need to elevate her permissions for this new role. We can go to Properties, Member Of, and then Add. Type 'domain admins' and hit OK. Then you can disable the administrative account and just use this account, which won't be well known to people outside of your system.

If you do disable the administrative account, be aware that it's the original admin account, and it's the only one that you can't lock with account lockout. So be careful about disabling it. You don't want to get locked out of your system or your server.

You'll also want to change passwords. Windows Server requires you to enter a complex password right after you've installed the operating system. You'll need to change any default account names and passwords and use complex passwords. To change your password, you just right-click on the account and go to Reset Password. That's relatively simple.

You'll also notice that there are a lot of built-in groups in Active Directory. The Users container, on the left, contains all of the built-in groups that come on your server by default. These groups can help you control what users can access and what they can do. For example, when you add somebody to the Backup Operators group, that user has the power to back up and restore files.

Let's go ahead and go to our Group Policy in Group Policy Manager. Go to Tools, and then we'll go to Group Policy Management. We're going to go to Domain Controllers and the Domain Controller default policy. Right-click. Edit. These are all the Group Policies that affect our domain controller. Over here, under Computer Configuration, we're going to go to Policies > Windows Settings > Security Settings.

You can see settings that you've enabled to lock down the domain controller behavior. We're going to look under Security Settings. We're going to look at Local Policies, specifically User Rights Management. Under here, on the right-hand side, we're going to go to Back Up Files and Directories, and you can see that right here. This is basically the built-in groups that can back up files and do backups in the directories.

We have administrators, backup operators, and server operators. We can add users to this group if we want to. Let's also go down to Restore Files and Directories. Here are the same groups. They all have that same access. All three of these, administrators, backup operators, and server operators, can back up and restore files on your system. That's something to be aware of. If I give someone Backup Operator rights, they'll be able to restore and back up any files on the system.

On the left side, there are also security options. For security options, we don't care who's logged on. They affect everyone. Scroll down to Interactive Logon. Let's do Interactive Logon. Don't require Ctrl + Alt + Del right here. This option reduces your security posture because the malware can take advantage of the fact that the Ctrl + Alt + Del box isn't there.

Having Ctrl + Alt + Del can prevent some malware. If this is enabled, it could open your system to some problems. But you can only set this across the board for every user. You can't define which user can do this; this setting applies to everyone. In this case, we don't want to do this. We're going to hit Cancel. Notice that it didn't ask us which user or which group to apply this to, since it affects everyone.

Other hardening strategies include installing only required software on the systems. With Windows Server, we have the concept of roles. Server Manager is quite a bit different as well. Here, we're going to go to Add Roles and Features. We can also see Remove Roles and Features. It has a wizard to get you started.

We'll just click through a few. This is very different from older versions of Windows Server, but you can do role-based or feature-based installation, or Remote Desktop Service installation. We're just going to do role-based. Select the server you want it to be on. I'm going to select my server. You can see the server roles that you can add. If we go on one more, you can see features.

Windows Server only installs binaries for any service that you want to run on the server. This helps reduce the attack surface. You won't have a bunch of programs open by default. You have to go in, enable them, and turn them on. One major difference between older servers and newer servers is that the security configuration wizard has been completely removed, so features are secured by default.

You should use Group Policy Manager or Microsoft Security Compliance Manager to control specific security settings.
We can also go to our services on Server Manager by selecting Tools and then Services.
It's a good practice to disable or turn off at least the startup behavior of any services that you're not requiring on this particular machine. Microsoft will start up services related to the roles you've installed, but often, there may be extra services that aren't required for the operation of your system. You can go through and look at all the services that are started automatically. For example, there's a service called SSDP that we're going to look at.
SSDP Discovery is a universal plug-and-play type of hardware on your system. It's universal plug-and-play for network devices. At one time, there were some flaws in it, and it's automatically disabled by default. You'll want to be careful not to disable a service that may be critical. Always do your research first.
That's it for this demo. In this demo, we discussed some strategies you can use to harden your Windows Server operating system.

8.1.4 Managing Automatic Updates (Demo Video)

Transcript:

In this demonstration, we're going to talk about configuring Windows Update.

As you probably know, every now and then, the Windows operating system goes out to Microsoft's site over the internet and checks whether there are any updates available for the operating system.
To manage updates on a Windows system, you need to come down to your Windows icon. Go to Settings and then Windows Update.
On the Windows Update screen, I can see the current status of my updates. Right here, you can see when my device was last updated. To be sure my system is up to date, I can go here and select Check for updates. When I do that, Windows tells me it's up to date.
Under Advanced Options, we next have the option to pause Windows updates for seven days. When I select this option, it'll pause my updates. And up here, you can see when the updates will resume. I don't want to pause my updates, so I'll click on Resume updates.
This option allows us to set the active hours on our system. A long time ago, when dinosaurs roamed the earth, there was nothing worse than working on a term paper and having Windows spontaneously decide it was a great time to install an update. The system would log you off, and hopefully, you'd have a saved copy of your work. If not, you might have to start all over.
So, this option lets you set active hours to let the device know when you typically work on your system. This is to avoid updating when you're in the middle of something without checking with you first. I like to work a lot, so I'll put in 16 hours. I'd like to change this to 7:00 a.m. until about 11:00 p.m.
Here, I can select Update History. When I do, this Window appears, and I can see the name of the updates and the date they were installed. If I need to know more about what an update is, I can click one of the links, and it takes me to the Microsoft Knowledge Base site for that particular update.
Below there are two options to be aware of. One is Uninstall updates, and the other is Recovery options. The Uninstall updates open another window where you can uninstall certain updates. However, some updates can't be removed from Windows.
Recovery options take you to where you can do a system reset, which isn't part of this demo.
Let's go back a few clicks.
Under Advanced Options, we do other things, like tell Windows to give us updates for other Microsoft products when we update. For example, I might want Office to be updated when updates run.
Here, we can tell Windows not to download when connected to metered connections. If we're connected to a cellular network and paying for bandwidth, we might not want to do that. Right now, that's off, and I'm going to leave it that way.
Here, you can tell Windows to restart as soon as possible when a restart is required to install an update. If you have this turned on, it'll warn you before restarting. You can tell Windows to show a notification when your PC requires a restart to finish an update by turning this On.
Okay, now let's move on to our delivery options. We're going to talk about a more advanced topic for a moment. It's possible to set up an update server on your local network and configure your workstations to get your updates from that server. This is a good option if you have a lot of workstations.
Instead of having all these workstations in your organization download the same files redundantly from Microsoft, wasting bandwidth, we'd want to download those updates once to the update server. Then all the systems on your

network would get the updates from that local server. Windows is taking that concept and making it more of a peer-to-peer model instead of a client-server model.
Essentially, what we can do is get our updates from Microsoft directly. We can also get updates from other workstations on our network that have already downloaded that same update. We can turn that functionality off or on right here. By default, it's turned On.
One option is only the PCs on my local network, which is selected by default. That's a pretty good idea. If you're a little more trusting, you can also get updates from computers out on the internet. This option scares me since I have no idea what I'd be getting from other computers, so I don't usually select it.
Right below this, we have Advanced options. Here, we can change the default settings for how much bandwidth we'll allow Windows to use for downloading and uploading updates on the system. We'll click back a few times, which wraps up Windows Update.
One last topic we need to cover before we end this demonstration: how to include device updates with Windows Update. This is usually a very good idea. The ironic thing is that finding this option can be challenging because when Windows releases featured updates, they often make things harder and harder to locate. That's what happened to the device update options.
First, we need to go to Bluetooth & devices settings. Now, go to Devices, More devices, and printer settings. I'm going to right-click the PC icon. From the menu, I'm going to select Device installation settings on this system.
Here, it says, Do you want to automatically download manufacturers' apps and custom icons available for your devices? This is set to Yes by default. I wanted to confirm that, so I'll leave it alone.
That's it for this demonstration. In this demo, we explained the process of managing Windows Update.

8.1.5 Configure Automatic Updates (Simulation)

Scenario

You need to customize how Windows Update checks for and installs updates on the ITAdmin desktop system.

In this lab, your task is to:

- Configure Windows Update to:
  - Install updates for other Microsoft products when Windows is updated.
  - Allow the installation of feature updates to be deferred 60 days.
  - Allow quality updates to be deferred 30 days.
- Configure Windows to automatically download manufacturers' apps and custom icons for devices.

Explanation

Complete this lab as follows:

- Configure the Windows Update settings.
  - Right-click **Start** and then select **Settings** .
  - Maximize the window for better viewing.
  - Select **Update & Security** .
  - From the right pane, select **Advanced options** .

Under Update Options, turn on **Receive updates for other Microsoft products when you update Windows** by sliding the switch to **On** .

Under *Choose when updates are installed* , configure each option as follows:

A feature update includes new capabilities and improvements. It can be deferred for **60** days.

A quality update includes security improvements. It can be deferred for this many days: **30**

Configure Windows to automatically download the manufacturer's apps and custom icons.

In the upper left of the Settings app, select the ***Home*** icon to return to the home page.

Select **System** .

From the left pane, select **About** .

From the right, under *Related settings* , select **Advanced system settings** .

Select the **Hardware** tab.

Select **Device Installation Settings** .

Select **Yes** and then select **Save Changes** .

Select **OK** .

8.1.6 Configure Microsoft Defender Firewall (Demo Video)

Transcript:

In this demonstration, we'll practice working with the Windows Defender Firewall. Windows Defender Firewall is a host-based firewall. It's implemented as software and designed to protect an individual system.

Be aware that Windows Defender Firewall was previously referred to as Windows Firewall. If you're on an older operating system, it may go by that name. However, it works basically the same. Its job is to prevent someone on the internet or the network from initiating an unwanted connection with a system. It just shuts it down.
To manage Windows Defender Firewall, I'll click on Search and type firewall, then select Windows Defender Firewall under Best match. You can see that the Windows Defender Firewall is currently Off. Let's click Turn Windows Defender Firewall on or off on the left. You'll see two network profiles to which the Windows Defender firewall is applied. We have the Private network settings and the Public network settings.
The Private network settings can be a little looser because we're assuming the company has a network firewall that's already blocking many connections, increasing the security level. We're trying to protect this individual host.
We want to customize the firewall for those times when we connect to a public network. Public networks, such as those defined at a restaurant, a hotel, or an airport, are like the wild, wild West of information technology. You have no idea who's on the network or their intentions. Therefore, you want to configure the Windows Defender Firewall with more stringent settings than when you're on a private network.
For both our network profiles, you can see that the Windows Defender Firewall is currently turned off. We need that firewall on. Select on for both. The default settings work well for a private network. We can see that we'll be notified if

Windows Defender Firewall blocks an app trying to communicate through the firewall. We'll have a list of apps allowed to communicate through the firewall.
That's probably appropriate. We'll run certain applications that need to send data back and forth through the host-based firewall. So that's okay. However, if I'm on a public network at an airport, a hotel, or a similar place, we want to turn this option on and block all incoming connections, including those in the list of allowed apps.
We don't want anyone on the public network initiating a connection with this computer, even if it's an allowed application. We'll leave this option checked for public network profiles and click OK. The Windows Defender Firewall state is on for incoming connections.
Let's talk about allowing certain apps. There may be times when you install an application that legitimately needs to communicate through the Windows Defender Firewall with other hosts on the network. Still, by default, it's going to be blocked. In that case, you can use the Allow an app or feature through Windows Defender Firewall. Let's select that.
You can see a list of applications and features that the Windows system is aware of. You can also see whether they're allowed through the firewall. For example, scroll down to Remote Assistance. Remote Assistance traffic is allowed through the host firewall for private and public profiles.
If there's an application that you need to allow through the firewall, you can enable it using the checkbox. For example, we've enabled Remote Desktop for Public and Private profiles. It's an app we frequently use for testing all sorts of networks.
There may be situations where you want to allow an application to communicate through the Windows Defender Firewall. Still, you can't find it in the list because Windows hasn't recognized it.
If that's the case, you can come down here and click Allow another app. Then, you locate the app. We have an app in the Program Files directory named VNC Viewer. VNC Viewer allows you to connect to desktops remotely for Windows, Linux, and other systems that won't work with a Remote Desktop.
We'll go ahead and select vncviewer. Click Add. Now you can see it's added. We can configure it to allow traffic from both the Public and Private network profiles. That's how you add an app to the list of apps that can communicate through the firewall. Click OK. Now, we've configured the Windows Defender Firewall on the system.
That's it for this demonstration. In this demo, we talked about managing the Windows Defender Firewall. We first turned on the Windows Defender Firewall and increased the security level of the Windows Defender Firewall for the public profile. We ended this demo by adding an exception to the firewall to allow a specific application to communicate through it.

8.1.7 Configure Microsoft Defender Firewall (Simulation)

Scenario

You have a new laptop that is running Windows 10. You notice a security message that indicates that Windows Firewall has been disabled. The laptop is currently connected to your organization's network, and the Domain network profile settings are in effect. You plan to travel this week and connect the laptop to various airport Wi-Fi hotspots. You need to enable Windows Firewall for any public network.

In this lab, your task is to configure Windows Firewall as follows:

- Turn on the Windows Firewall for the Public network profile.
- In addition to the programs and ports currently allowed, allow the following service and programs through the firewall for only the Public network profile:
  - A service named **Key Management Service**
  - An application named **Arch98**
  - An application named **Apconf**

Explanation

To complete this lab, you need to allow the following service and programs through the firewall for the Public network profile only:

A service named **Key Management Service**

An application named **Arch98**

An application named **Apconf**

Leave all other existing apps and features as they are.

Complete this lab as follows:

Turn on the firewall for the public network.

Right-click **Start** and then select **Settings** .

Select **Network & Internet** .

From the right pane, scroll down and select **Windows Firewall** .

From the *Firewall & network protection* dialog, under Public network, select **Turn on** .

Allow applications to communicate through the firewall for the Public network only.

Select **Allow an app through firewall** .

Select **Change settings** .

For *Key Management Service,* clear **Domain** and **Private** , and then select **Public** .

Select **Allow another app** to configure an exception for an application not currently allowed through the firewall.

Select the ***application*** from the list and then select **Add** .

For the newly added application, clear **Domain** and **Private** , and then select **Public** .

Repeat steps 3d - 3f for the remaining ***application*** .

Select **OK** .

8.1.8 Practice Questions (Section Quiz)

q_harden_allow_block_lists_secp8

You are the IT Security Specialist at a financial institution. The company has just set up a new server that will host a critical application.

As part of the company's OS hardening policy, you are tasked with ensuring the server is secure before it goes live. One of the key requirements is to minimize the server's attack surface.

What is the MOST effective way to achieve this?

Answers:

Install the latest antivirus software on the server.

***Implement allow and block lists to control which applications and services can run on the server.**

Regularly update the server's operating system and all installed applications.

Enforce a strict password policy for all users who will have access to the server.

Explanation:

Implementing allow lists and block lists is the correct answer. It is an effective way to control which applications and services can run on the server. By only allowing necessary applications and services to run, you can significantly reduce the server's attack surface.

While installing the latest antivirus software is an important step in securing the server, it does not directly contribute to minimizing the server's attack surface. The attack surface is reduced by limiting the number of applications and services that can run on the server.

Regularly updating the server's operating system and all installed applications is a crucial part of maintaining security, but it does not directly minimize the server's attack surface. The attack surface is minimized by limiting the number of applications and services that can run on the server.

Enforcing a strict password policy is an important part of securing the server, but it does not directly contribute to minimizing the server's attack surface. The attack surface is reduced by limiting the number of applications and services that can run on the server.

q_harden_auto_updt_secp8

Which of the following tools can you use on a Windows network to automatically distribute and install software and operating system patches on workstations? (Select two.)

Answers:

***Group Policy**

***WSUS**

Security Templates

Security Configuration and Analysis

Configuration baseline

Explanation:

Windows Software Update Services (WSUS) is a patch management tool that allows clients on a network to download software updates from an internal WSUS server in their organization.

The WSUS server receives a list of available updates from Microsoft.

On the WSUS server, you identify allowed or required patches for your organization.

Clients download only approved patches from an internal WSUS server or directly from Microsoft.

You can also use Group Policy to distribute and automatically install patches. You must use Group Policy to install updates to non-Microsoft software that is not supported with WSUS.

Use the Security Templates snap-in to create and edit templates that enforce system security settings.

Use the Security Configuration and Analysis snap-in to compare the existing settings with the template or to apply a template to a single device.

A configuration baseline is a set of consistent requirements for a workstation or server. However, baselines are not directly associated with automatically distributing and installing software and operating system patches on workstations.

q_harden_baseline_secp8

Which of the following BEST describes a configuration baseline?

Answers:

***A list of common security settings that a group or all devices share.**

A collection of security settings that can be automatically applied to a device.

A set of performance statistics that identifies normal operating performance.

The minimum services required for a server to function.

Explanation:

A configuration baseline is a set of consistent requirements for a workstation or server. Configuration baselines include a component that ensures that all workstations and servers comply with the security goals of the organization.

A security template is a saved set of configuration values that produce the system configuration specified in the configuration baseline. When you apply the security template to a system, the settings within the template are applied to the system.

A performance baseline is a set of performance statistics that identify normal operating performance.

q_harden_configuration_enforcement_secp8

You are the IT Security Manager at a large corporation. The company has just acquired a batch of new servers that will be used to host critical applications.

As part of the company's OS hardening policy, you are tasked with ensuring these servers are secure before they are deployed.

What should be your first step?

Answers:

Install the latest antivirus software on the servers.

Enforce a strict password policy for all users who will have access to the servers.

***Apply a configuration baseline template to the servers to ensure they are configured according to industry security standards.**

Connect the servers to the network to monitor for any potential security threats.

Explanation:

Applying a configuration baseline template is the correct answer, as it ensures that the servers are configured according to industry security standards. This reduces the potential attack surface and is the first step in securing the servers.

While installing the latest antivirus software is an important step in securing the servers, it should not be the first step. Before any software is installed, the servers should be configured to a secure baseline.

Enforcing a strict password policy is also an important part of securing the servers, but it is not the first step. Before users are given access to the servers, they should be configured securely.

Connecting the servers to the network before they have been properly secured could expose them to potential security threats. This should only be done after the servers have been configured securely.

q_harden_harden_secp8

By definition, what is the process of reducing security exposure and tightening security controls?

Answers:

Social engineering

***Hardening**

Active scanning

Passive reconnaissance

Explanation:

Hardening is the process of securing devices and software by reducing security exposure and tightening security controls.

Social engineering is the act of exploiting human nature by convincing someone to reveal information or perform an activity.

Active scanning and passive reconnaissance are types of reconnaissance attacks.

q_harden_limit_privileges_secp8

A small IT company is experiencing an increased number of cyberattacks. Its server uses default settings from the developer, which the company believes is a potential source of vulnerability.

Which of the following changes should the IT company consider to improve the server's security?

Answers:

***Limit the privileges of each user on the server.**

Continue to use the default settings but increase monitoring efforts.

Switch off the server whenever it is not in use.

Ignore software and security patches and updates.

Explanation:

Limiting the privileges of each user on the server to the least amount necessary to perform their function reduces the impact of a compromised account. This is part of the principle of least privilege and is a recognized method for hardening server security.

Although increasing monitoring efforts can help identify and respond to security incidents, it does not fundamentally reduce the potential vulnerabilities in the default server settings.

Switching off the server when not in use may reduce the exposure to potential attacks, but it may not be feasible due to the need for constant availability.

Ignoring software security patches and updates is risky behavior; the IT department should apply them consistently and timely as part of a secure configuration approach.

q_harden_removable_media_encryption_secp8

You are the IT Security Manager at a healthcare organization. The organization frequently uses removable media devices like USB drives to transfer patient data between different departments.

Given the sensitive nature of the data, you are tasked with ensuring the data remains secure during transfer.

What is the MOST effective way to achieve this?

Answers:

Implement a strict password policy for all users who will have access to the data.

Install the latest antivirus software on all computers that will be used to transfer the data.

***Encrypt the data on the removable media devices.**

Regularly update the operating systems and all installed applications on the computers that will be used to transfer the data.

Explanation:

Encrypting the data on the removable media devices is the correct answer. It ensures that even if the devices are lost or stolen, the data cannot be accessed without the encryption key.

While implementing a strict password policy is an important part of securing the data, it does not directly protect the data during transfer. The data is most vulnerable when it is stored on the removable media devices, so it should be encrypted to ensure its security.

Installing the latest antivirus software is an important step in securing the computers, but it does not directly protect the data during transfer. The data is most vulnerable when it is stored on the removable media devices, so it should be encrypted to ensure its security.

Regularly updating the operating systems and all installed applications is a crucial part of maintaining security, but it does not directly protect the data during transfer. The data is most vulnerable when it is stored on the removable media devices, so it should be encrypted to ensure its security.

q_harden_role_separation_secp8

As a system administrator, you are tasked with hardening the system in your organization.

Which of the following strategies is MOST effective in reducing the impact if a single system is compromised?

Answers:

Installing anti-virus software on all systems

Regularly updating all software to the latest version

***Role separation**

Using complex passwords

Explanation:

Role separation is the correct answer. Role separation involves installing services on separate physical systems. This means that if a single system is compromised, only the few services on that system will be affected. This strategy is most effective in reducing the impact if a single system is compromised.

While installing anti-virus software on all systems is a good security practice, it does not specifically address the issue of reducing the impact if a single system is compromised. Anti-virus software helps to protect individual systems from malware, but it does not prevent the spread of a compromise from one system to another.

Regularly updating all software to the latest version is also a good security practice, as it ensures that all known vulnerabilities are patched. However, this does not specifically address the issue of reducing the impact if a single system is compromised.

Using complex passwords is a good security practice, as it makes it harder for unauthorized users to gain access to systems. However, this does not specifically address the issue of reducing the impact if a single system is compromised.

q_harden_security_baseline_secp8

A global corporation has faced numerous cyber threats and is now prioritizing the security of its servers. The corporation's IT security expert recommends a strategy to improve server security.

Which of the following options is likely to be the MOST effective?

Answers:

***Implement a security baseline, consistently apply updates and patches, and adhere to hardening guidelines.**

Enable all services on the servers to maximize functionality.

Utilize easily-remembered, simple passwords to improve server manageability.

Switch off all security features to enhance server performance and reduce latency.

Explanation:

By establishing a security baseline, applying regular updates and patches, and adhering to hardening guidelines, the servers' security can be significantly improved, reducing the risk of cyber threats.

Enabling all services on a server may unnecessarily increase its attack surface. This strategy is more likely to make servers vulnerable to cyber attacks.

Despite the ability to easily remember them, simple passwords are a significant security risk as cyber attackers can easily guess or crack them.

Turning off all security features might enhance performance but significantly compromise server security. It is crucial to strike a balance between performance and security, and under no circumstances should the IT security expert disable all security features.

q_harden_services_secp8

Which of the following actions should you take to reduce the attack surface of a server?

Answers:

***Remove unused services.**

Install anti-malware software.

Install the latest patches and hotfixes.

Install a host-based IDS.

Explanation:

Attack surface reduction (ASR) cuts down on the software or services running on a system. By removing unnecessary software, features, or services, you eliminate possible attacks directed at those components. When removing unnecessary components, you should:

Use role separation by installing services on separate physical systems. If a single system is compromised, only the few services on that system are affected.

For many new systems, unnecessary services are often installed by default. Following installation, remove unneeded services, protocols, and applications.

When removing existing services, determine the unneeded services and their dependencies before altering the system.

Adding anti-malware or a host-based intrusion detection system (IDS) adds a level of protection (defense in depth) but does not reduce the number of components running on the system.

Applying patches is necessary to fix security problems with software or the operating system. But if the system is not running a specific piece of software, the patches that apply to that software are irrelevant and do not need to be applied.

q_harden_settings_secp8

Which of the following do security templates allow you to do? (Select two.)

Answers:

***Quickly apply settings to multiple computers.**

Apply new software patches.

***Configure consistent security settings between devices.**

Block malicious websites.

Fix a specific software problem.

Explanation:

Security templates allow you to quickly and consistently apply settings to multiple computers in order to bring them into compliance with a security baseline.

Security templates are not used to apply new patches, block malicious websites, or fix specific software problems.

q_harden_test_patches_secp8

You are the IT manager of a mid-sized company. Your team has just received a critical security patch for your company's primary database software. The patch is meant to address a significant vulnerability that could potentially expose sensitive customer data.

However, applying the patch requires taking the database offline, which would disrupt operations.

What should you do?

Answers:

Apply the patch immediately to all systems to ensure the vulnerability is addressed as soon as possible.

Ignore the patch. The disruption to operations is too significant.

***Apply the patch to a test environment first, then roll it out gradually to the rest of the systems.**

Wait for a few weeks before applying the patch to see if any issues arise from other users.

Explanation:

Apply the patch to a test environment first is the correct answer. Testing the patch in a controlled environment allows you to identify and address any issues before rolling it out to the rest of the systems. This approach minimizes the risk of disruption to operations.

Applying the patch immediately to all systems could potentially cause unforeseen issues. If the patch has any bugs or compatibility issues, it could cause more harm than good.

Ignoring the patch is not advisable because it leaves the company vulnerable to the security threat the patch is meant to address. This could lead to a data breach, which would be far more disruptive to operations than applying the patch.

Waiting for a few weeks before applying the patch is not a good strategy. The longer the vulnerability remains unaddressed, the higher the risk of a security breach. It's important to address vulnerabilities as soon as possible, but in a controlled and measured way.

q_harden_third_party_secp8

You have just purchased a new network device and are getting ready to connect it to your network.

Which of the following actions should you take to increase its security? (Select two.)

Answers:

***Change default account passwords.**

Remove any backdoors.

***Apply all patches and updates.**

Conduct privilege escalation.

Implement separation of duties.

Explanation:

To secure new devices, apply all recent patches and updates and then change the default user account passwords. For some systems, you can also increase security by changing the default account usernames. Default account usernames and passwords are well-known and can be easily discovered.

A backdoor is an unprotected access method or pathway. Backdoors are added by attackers or programmers during development. Backdoors that are present on new devices are typically hard-coded and must be removed by editing the code.

Privilege escalation allows a user to take advantage of a software bug or design flaw in an application to gain access to system resources or additional privileges that are typically not available to normal users. Separation of duties is the concept of requiring the participation of at least two people to complete a task. This helps prevent insider attacks because no one person has end-to-end control, and no one person is irreplaceable.

8.2 File Server Security

As you study this section, answer the following questions:

How can you identify inherited permissions?

How do Share and NTFS permissions differ?

On which elements can NTFS permissions be set?

How can you view the users who have permissions for a particular drive?

In this section, you will learn to:

- Configure NTFS permissions.
- Disable inheritance.

Key terms for this section include the following:

Term	Definition
Shared folder	A folder whose contents are available over the network.
Network-attached storage (NAS)	A standalone storage device or appliance that acts as a file server.
Storage area network (SAN)	A special network composed of high-speed storage that is shared by multiple servers.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts. Physical security 1.4 Explain the importance of using appropriate cryptographic solutions. Encryption Level Full-disk File 2.2 Explain common threat vectors and attack surfaces. File-based 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Access control

Access control list (ACL)

Permissions

Least privilege

Hardening techniques

Removal of unnecessary software

3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

Secure communication/access

Virtual private network (VPN)

Tunneling

Internet protocol security (IPSec)

4.1 Given a scenario, apply common security techniques to computing resources.

Hardening targets

Servers

4.2 Explain the security implications of proper hardware, software, and data asset management.

Acquisition/procurement process

4.3 Explain various activities associated with vulnerability management.

Identification methods

System/process audit

4.5 Given a scenario, modify enterprise capabilities to enhance security.

Implementation of secure protocols

File integrity monitoring

4.6 Given a scenario, implement and maintain identity and access management.

Provisioning/de-provisioning user accounts

Access controls

	Least privilege
TestOut Security Pro	3.1 Harden computer systems 3.1.1 Configure file system inheritance 3.1.3 Configure NTFS permissions 4.2 Implement Encryption Technologies 4.2.2 Encrypt files

8.2.1 File Server Security (Lesson Video)

Transcript:

File servers are an attractive target because they almost always store some type of valuable data, such as client contacts, credit card information, employee files, company performance records--the list goes on and on. In this lesson, we'll look at some steps you should take to properly secure your organization's file servers.

Let's start by looking at some of the best practices you should follow when it comes to file server security. Keep in mind that these are generic best practices that can be applied to any type of file server, whether it's a dedicated NAS, a SAN, a Microsoft server, or a Linux server.

The first best practice is to keep all file servers physically secure. Remember, it's very easy for someone to gain access to data if they have physical access to a system, so employ the appropriate physical controls with your file servers as you would with any other networking hardware device.

Next, is the principle of least privilege. This is the concept that users should only have access to the information and resources that are necessary for them to perform their job, and that's it. They shouldn't have any more privileges than they need. This helps prevent privilege creep, where a user slowly accumulates more and more privilege than they need to perform their job.

Another best practice is to use full-disk encryption on all storage mediums. This may or may not be possible for your organization depending on its policies. Encrypting an entire drive and keeping it encrypted as people work requires a significant amount of resources. Due to this resource overhead, it might not be feasible to use full-disk encryption. Instead, you might want to consider only encrypting backups.

Another option is to fully decrypt drives when they're powered on and encrypt them when the system shuts down. This may not prevent someone from transferring files to a USB drive, but does prevent someone physically stealing the entire drive or server.

Another security best practice that applies to almost every system, not just files servers, is to remove unnecessary software and disable unused services. You can think of this as a principle of least privilege with applications.

In other words, you should only have the applications and services that are necessary for the file server to function, and nothing more. Having unused applications or services increases the attack surface area of a system and makes securing it that much harder. In addition, the system is more susceptible to zero-day attacks.

The next practice you should employ is auditing. Make sure that your file server has proper auditing enabled that can track user behavior, such as when files are accessed, modified, deleted, and moved. This helps you identify abnormal behavior on the system and can even help you detect that your system has been compromised.

The last security best practice is to use implicit deny access control lists, or ACLs. ACLs control the system's permission. A good security practice is to prevent access to everything that isn't explicitly granted. In other words, deny access to everyone except people with explicitly granted permissions. This control is a part of the concept of least privilege mentioned earlier.

The opposite of implicit deny is explicit deny, meaning that you're only looking for people you want to deny access to. The problem with this approach is that you have to do a bit of guesswork and research to deny access to files. And if you happen to forget about a particular group or file tree, you could inadvertently give everyone access to sensitive files. As a result, your system will be more secure if you use implicit deny.

At some point, we all need to download or transmit files using the internet, which is inherently unsecure. Let's talk about different ways files can be transmitted across a network. Some of these protocols are secure, and some aren't. As a security administrator, you need to know the difference and decide which of these protocols to use on your systems.

File Transfer Protocol, or FTP, is an older TCP/IP protocol that's used for transferring files across systems. It was adopted as a TCP/IP standard through RFC 959 in 1985. FTP is inherently unsecure since it doesn't encrypt data, including usernames and passwords. As such, FTP shouldn't be used for transferring sensitive data over an unsecure network such as the internet. Another problem with FTP is that it allows anonymous user access. This option can be disabled, but it's still a security risk. If you must use FTP, there are ways to use it more securely by pairing it with an encryption protocol. There are a couple of ways to do this.

The first way is to use a VPN connection that wraps all communications, FTP packets included, in an encrypted data packet. This would encrypt all traffic between the host and the FTP file server. IPsec or SSH can be used to create secure tunneling with FTP to secure the transmission. When use with an SSH tunnel, FTP is called FTP over SSH, or Secure FTP.

Another way is to use FTPS. FTPS uses SSL and TLS encryption to secure data transmissions. FTPS is similar to HTTPS in that it uses the SSL and TLS cryptographic protocols to secure communications and also requires the use of SSL certificates to encrypt traffic.

Another protocol that you can use is the Secure Copy Protocol, or SCP. SCP is similar to the Linux copy command cp. SCP uses the Secure Shell Protocol, or SSH, which is a secure tunnel. SCP is an older protocol that's entirely text-basedâ€”that is, you need to type the command, including the directories to transmit, into a command line interface. While it's secure, it doesn't offer any type of directory traversing.

As a replacement for SCP, you can use SFTP. SFTP stands for Secure Shell File Transfer Protocol. Files are transferred through an SSH tunnel. SFTP isn't related to FTP at all. They're two completely different, incompatible protocols. SFTP typically uses SSH v2, the more secure version of Secure Shell Protocol. It uses a graphical interface, and it's more user-friendly.

That's it for this lesson. In this video, we discussed several best practices that can be used to secure your file servers and your organization's resources. We also discussed how to use an encrypted protocol to reduce the opportunity for an attacker to see sensitive information as you copy files across a network.

8.2.2 File System Security Facts

Managing the file system is a primary concern of IT professionals. The file system is responsible for storing and securing data. An organization depends on data and requires that it be secure and easily accessible.

This lesson covers the following topics.

- File system security
- Least privilege permission assignments
- Data transfer security protocols

File System Security

Tasks to secure file servers include:

- Prevent physical access.
- Implement the principle of least privilege.
- Use full-disk encryption on backups.
- Implement strong authentication.
- Remove unnecessary software and disable unused services.
- Use implicit deny access control lists (ACLs).

Use hidden folders and files.

Audit the file system.

When managing the security of the file system, be aware of the following:

The transfer of files between a client and a server is often unsecured. Use IPSec or a VPN between the server and the client to secure data as it travels through the network.

File and print resources are primarily vulnerable to denial-of-service (DoS) and access attacks.

Attacks on file servers are often directed against the NetBIOS protocol. To protect the server, verify that NetBIOS is not required on the server, disable the NetBIOS protocol on the server, and use a host-based firewall to close NetBIOS ports 135 and 137 - 139.

A *shared folder* is a folder whose contents are available over the network.

An *administrative* share is a shared folder that is available only to an administrative user.

Administrative shares are hidden. This means that the share will not display when a user browses a network computer.

By default, the root of every drive is an administrative share.

In Windows, you can create hidden shares by appending a $ sign to the end of the share name (for example, DataFiles$).

Users must know the name of the share to access it and have the appropriate access permissions.

Do not share the root directory with regular users.

With Windows Server 2008 and later, you can use File Server Resource Manager (FSRM) to control files saved on a file server.

Quotas limit the amount of data that can be saved within a folder. A hard limit prevents exceeding the quota limit, while a soft limit sends a message when the limit is exceeded.

File screens restrict the type of files that can be saved in a folder. For example, you can prevent media files (audio and video) or files with specific file extensions from being saved. An *active* file screen prevents saving the specified file types, while a *passive* screen monitors when the specified file types are added to the folder.

Least Privilege Permission Assignments

Least privilege means that a principal is granted the minimum possible sufficient rights to complete a task they are authorized to perform. This mitigates risk if an account should be compromised and fall under the control of a threat actor. Least privilege involves a design phase, where analysis of business workflows determines what roles and permissions are required.

While least privilege is a strong design principle, implementing it successfully can be challenging. Where many users, groups, roles, and resources are involved, managing permission assignments and implications is complex and time-consuming. Improperly configured accounts can have two different impact s . On the one hand, setting privileges that are too restrictive

creates a large volume of support calls and reduces productivity. On the other hand, granting too many privileges to users weakens the system's security and increases the risk of malware infection and a data breach.

Ensuring least privilege also involves continual monitoring to prevent authorization creep. Authorization creep refers to a situation where a user acquires more and more rights, either directly or by being added to security groups or roles.

For example, a user may be granted elevated privileges temporarily. In this case, a system is needed to ensure that the privileges are revoked at the end of the agreed period. A system of auditing should regularly review privileges, monitor group membership, review access control lists for each resource, and identify and disable unnecessary accounts.

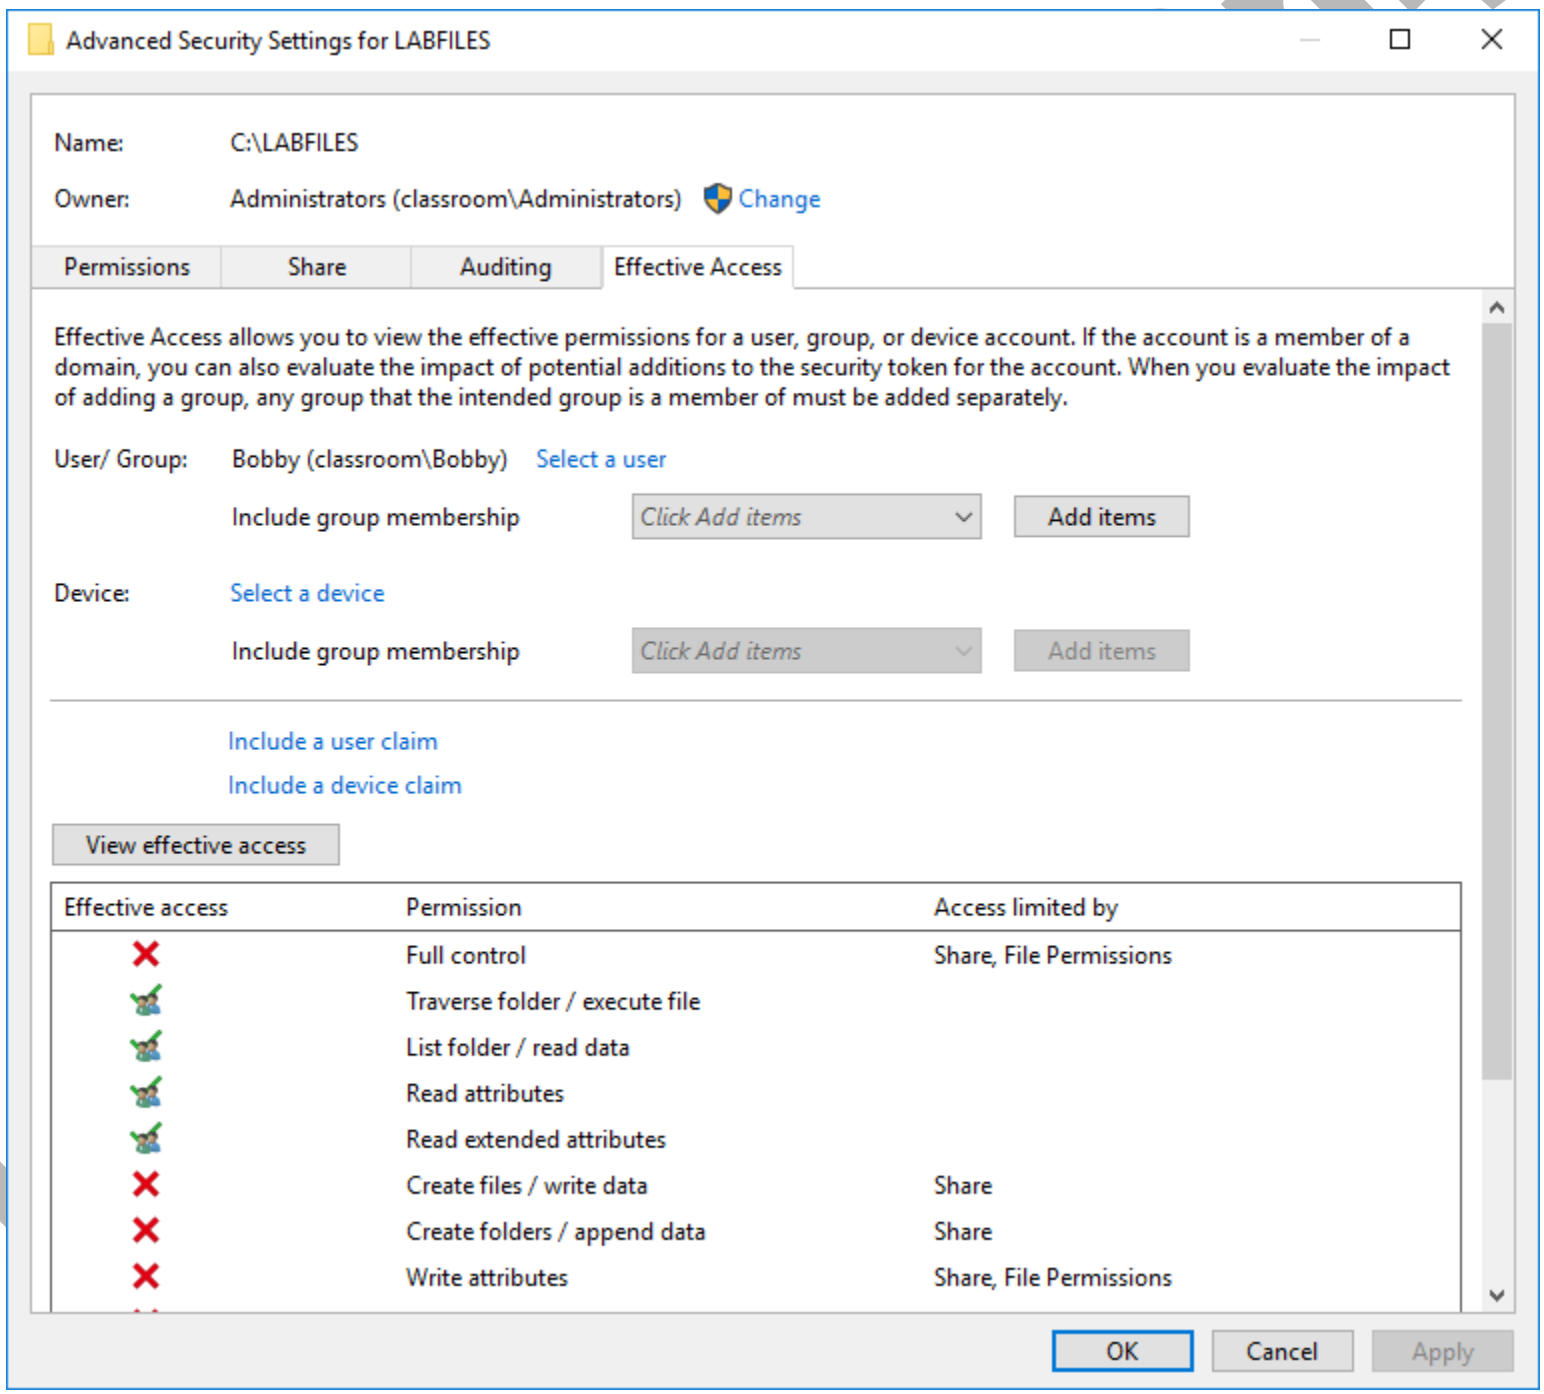

Determining effective permissions for a shared folder. (Screenshot used with permission from Microsoft.)

Provisioning is the process of setting up a service according to a standard procedure or best practice checklist. The IT department must keep track of all assets under management, and user accounts are a type of asset. User accounts are provisioned for new employees and temporary access, such as by consultants and contractors. Some businesses may also need to set up customer accounts.

Provisioning a user account involves the following general steps:

Identity Proofing - verifies that the person is who they say they are by checking official documents and records. Circumstances might also demand a background check, which verifies current and previous addresses, education, or previous employment and whether the person has a criminal record or credit issues.

Issuing Credentials - allows the user to select a password known only to them and/or enroll them with biometric or token-based authenticators.

Issuing Hardware and Software Assets - the user will typically need a computer or smartphone and possibly local copies of licensed software apps. Employees need sufficient resources to do their job. If their resources are inadequate, they might try to obtain hardware and software directly (shadow IT).

Teaching Policy Awareness - by scheduling training and providing access to learning resources so that the employee or contractor is aware of security policies and risks. They must also be aware of policies for the personal use of any IT assets issued to them.

Creating Permissions Assignment - by identifying the work roles that the account must support and configuring the appropriate rights using a role-based, mandatory, or attribute-based access control model. If the account is granted privileged access, it should be tagged for close monitoring.

Deprovisioning is the process of removing the access rights and permissions allocated to an employee when they leave the company or from a contractor when a project finishes. This involves removing the account from any roles or security groups. The account might be disabled for a period and then deleted or deleted immediately.

Data Transfer Security Protocols

The following table describes considerations for securing file transfer using TCP/IP protocols:

Protocol	Description
File Transfer Protocol (FTP)	Be aware of the following when using FTP: *Anonymous login* (also known as *blind* or *anonymous FTP*) allows unrestricted access to the FTP server. Disable anonymous login to control access based on username. The username and password are transferred in cleartext and can be captured in transit by a sniffer. To protect logon credentials, implement a secure protocol, such as Secure Socket Layer (SSL). Use IPSec or a VPN tunnel to protect data transfers.

	The write permission allows users to upload files to the FTP server. Carefully restrict which users have the write permission. FTP uses port 21 for control information (such as logon) and port 20 for data transfer.
Trivial File Transfer Protocol (TFTP)	TFTP provides no authentication, encryption, or error detection. In addition, TFTP uses UDP instead of TCP. TFTP might be faster than FTP, but it does not perform error detection which could result in file errors.
Secure Copy Protocol (SCP)	SCP uses Secure Shell version 1 (SSH1) to secure file transfers and login credentials.
Secure Shell File Transfer Protocol (SFTP)	SFTP is a file transfer protocol that uses Secure Shell version 2 (SSH2) to secure data transfers. SFTP is not an FTP that uses SSH but rather a secure transfer protocol that is different from FTP.
Secure FTP	Secure FTP (also known as FTP over SSH) tunnels FTP traffic through an SSH tunnel.
FTP Secure (FTPS)	FTPS adds SSL or Transport Layer Security (TLS) to FTP in order to secure logon credentials and encrypt data transfers. FTPS requires a server certificate.

8.2.3 File Permission Facts

This lesson covers the following topics:

- File system permissions
- Managing file system permissions

File System Permissions

With file system security, each object in the file system has an access control list (ACL) associated with it. The ACL contains a list of accounts (principals) allowed to access the resource and the permissions they have over it. The order of ACEs in the ACL is important in determining effective permissions for a given account. ACLs can be enforced by a file system that supports permissions, such as NTFS, ext3/ext4, or ZFS.

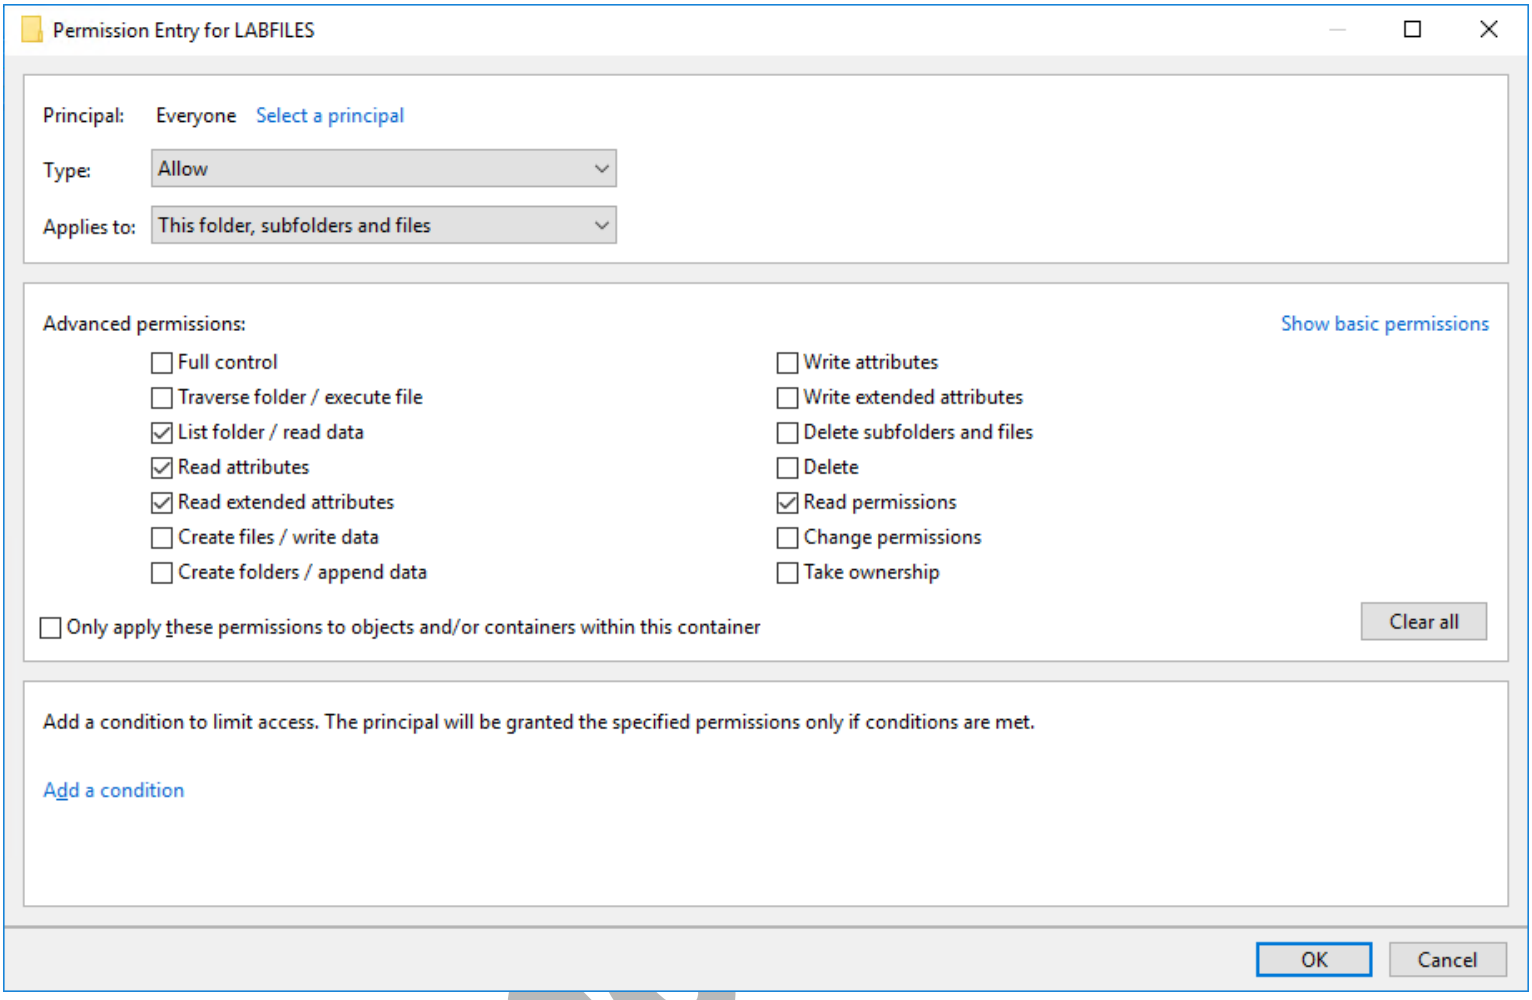

Configuring an access control entry for a folder. (Screenshot used with permission from Microsoft.)

For example, in Linux, there are three basic permissions:

Read (r) - is the ability to access and view the contents of a file or list the contents of a directory.

Write (w) - is the ability to save changes to a file or create, rename, and delete files in a directory (also requires execute).

Execute (x) - is the ability to run a script, program, or other software file, or the ability to access a directory, execute a file from that directory, or perform a task on that directory, such as file search.

These permissions can be applied in the context of the owner user (u), a group account (g), and all other users/world (o). A permission string lists the permissions granted in each of these contexts:

drwxr-xr-xhome

The string above shows that for the directory (d), the owner has read, write, and execute permissions, while the group context and other users have read and execute permissions.

The **chmod** command is used to modify permissions. It can be used in symbolic mode or absolute mode. In symbolic mode, the command works as follows:

chmod g+w, o-xhome

The effect of this command is to append write permission to the group context and remove execute permission from the other context. By contrast, the command can also be used to replace existing permissions. For example, the following command applies the configuration shown in the first permission string:

chmod u=rwx,g=rx,o=rxhome

In absolute mode, permissions are assigned using octal notation, where r=4, w=2, and x=1. For example, the following command has the same effect:

chmod 755home

In this example, the 755 correlates to the permissions assigned to the user, group, and others, where user permissions are represented by 7, group permissions are 5, and others are also 5. The numbers are generated by adding the values associated with r (read), w (write), and x (execute). The only combination of values that can result in 7 is 4+2+1 or r, w, x. Similarly, the only combination of values resulting in 5 is 4+1, or r,x. This means the owner has r,w, and x, whereas the group and others have only r and x.

Managing File System Permissions

On a Windows system, access to files is controlled through two sets of permissions: share and New Technology File System (NTFS). The following table describes permissions specific to each type.

Permission Type	Description
Share	Share permissions control access through a network connection with the file server. If files are accessed locally, share permissions do not control access. Share permissions have three levels of permissions: Reader (read-only) Contributor (read and write) Owner or Co-owner (full control or all permissions) Share permissions can be set only on a folder.
NTFS	NTFS permissions: Can be set on drives, folders, and files. Control both local and network access.

Permission Type	Description
	Have dozens of permissions that offer granular control over what actions are allowed. Can be set only on volumes formatted with NTFS.

Be aware that:

- Both share and NTFS permissions use a discretionary access control list (DACL) for controlling access. The access list identifies the users or groups and their associated permissions to files or folders.
- Both share and NTFS permissions include Allow or Deny permissions. Deny permissions override Allow permissions.
- Both share and NTFS permissions must be configured to allow access through the share. If a user is allowed share access but no NTFS permissions are set for the user or a group to which the user belongs, no access will be allowed.
- Effective permissions to shared folders are the more restrictive of either share or NTFS permissions.
- A user's effective permissions cannot be greater than the share permissions assigned to the user or a group to which the user belongs. For this reason, a common strategy for combining share and NTFS permissions is to:
  - Assign Co-owner share permissions to Everyone.
  - Use NTFS permissions to control access. Use the principle of least privilege by assigning NTFS permissions only to necessary groups and by assigning only the necessary permissions to those groups. Even though Everyone has share permissions, only the users or groups with NTFS permissions will have access.
- Permissions for folders and files can be inherited. On Windows systems, the Advanced Security settings identify when permission inheritance is in effect.
- Whenever possible, assign permissions to groups rather than users. Users receive the permissions assigned to their groups.

8.2.4 Configuring NTFS Permissions (Demo Video)

Transcript:

In this demonstration, we're going to work with NTFS and share permissions. Since these two permissions overlap, it's important to understand how they relate to each other. I'm on two different machines here. On one machine, I'm logged in as the TestOut user. On the other, I'm logged in as the jdoe user.

We're going to share from this machine. I'll go to File Explorer > This PC > C: > Research. Now I'll right-click and select Properties > Sharing. You can see that this folder is currently shared on the network. I'll click on the Share button. Currently, only the TestOut user has rights to access the share—both Read and Write access. Let's close this.
Let's go to our Security tab, which houses our NTFS file system permissions. Currently, the System group, the TestOut user, and the Administrators group have NTFS permissions to the Research folder. The jdoe user—which is the user currently logged in to the other system—doesn't have any access. I'll switch to the other machine and try to connect to

the share. Since I'm logged in as jdoe, those are the credentials that'll be used to try to connect to the Research folder. When I browse to it and click on it, it pops up and explains that I don't have permission to access this folder.
Let's go back to the first machine and change it so the jdoe user can access the network share. Again, on the Properties page for the Research folder, I need to add permissions for jdoe. I'll change the Permission Level to Read/Write and then click Share.
Now let's go to our NTFS permissions found on the Security tab and view the jdoe user's permissions. Notice that the jdoe user has been added to the NTFS permissions automatically when we added the user to the share. As you can see, jdoe has Full control over files on this share. So, both share and NTFS permissions have been given to the jdoe user for the Research folder. Let's test from the other machine. Browsing back to the share, instead of an error, the jdoe user has access now.
Now, what happens when your NTFS and share permissions aren't aligned or they conflict with each other? Let's change the settings and see what happens. Back on the Properties page of the Research folder, let's change the NTFS settings so the jdoe user doesn't have the ability to change files. We just click Edit, select the user, and then remove the Full control, Modify, and Write permissions. Then we click OK. Again, switching back to the other machine, I'll open the Research folder and open a file to make changes to it. When I click Save now, it asks me where to save the file because the current location—which is actually on the other machine—isn't allowing me to save changes to the file. I can only save a copy to a different location. So, the jdoe user has permission to change the file as far as the network share is concerned, but the NTFS permissions don't allow any changes.
If you go back in to share rights and give jdoe access to read the files, but you don't change the NTFS rights, you'll have the same problem. It's not that share rights or NTFS rights supersede the other. It's that whichever one is more restrictive will be enforced. So since the NTFS rights are more restrictive here, they supersede all the other rights.
Okay, that's it for this demonstration. In this demo, we talked about configuring share and NTFS permissions. We first looked at how to assign NTFS and share permissions for a folder and for a user account. And we talked about how rights are applied when there's a conflict between the two types of permissions.

8.2.5 Configure NTFS Permissions (Simulation)

Scenario

There are two groups of users who access the CorpFiles server: Marketing and Research.

Each group has a corresponding folder:

D:\Marketing Data

D:\Research Data

In this lab, your task is to:

Disable permissions inheritance for **D:\Marketing Data** and **D:\Research Data** and convert the existing permissions to explicit permissions.

For each of the above folders, remove the **Users** group from the access control list (ACL).

Add the **Marketing** group to the Marketing Data folder ACL.

Add the **Research** group to the Research Data folder ACL.

Assign the groups **Full Control** to their respective folders.

Do not change any other permissions assigned to other users or groups.

Explanation

While completing this lab, use the following:

D:\Marketing Data

D:\Research Data

Groups:

Marketing

Research

Complete this lab as follows:

Open the Data (D:) drive.

From the Windows taskbar, select **File Explorer** .

From the left pane, expand and select **This PC** > **Data (D:)** .

Disable inheritance and convert inherited permissions to explicit permissions.

From the right pane, right-click the applicable ***folder*** and then select **Properties** .

Select the **Security** tab.

Select **Advanced** to modify inherited permissions.

Select **Disable inheritance** to prevent inherited permissions.

Select **Convert inherited permissions into explicit permissions on this object** .

Remove the Users group from the access control list.

Under *Permission entries* , select **Users** .

Select **Remove** to remove the group from the access control list.

Select **OK** .

Add a new group to the access control list and allow Full Control.

From the Properties dialog, select **Edit** to add a group to the access control list.

Select **Add** .

Enter the ***name*** of the group you want to add and then select **Check Names** .

Select **OK** .

With the newly added group selected, under the Allow column, select **Full control** and then select **OK** .

Select **OK** to close the properties dialog.

Repeat steps 2 - 4 to modify the permissions for the additional folder.

8.2.6 Disable Inheritance (Simulation)

Scenario

Confidential personnel data is stored on the CorpFiles file server in a shared directory named Personnel. You need to configure NTFS permissions for this folder so that only managers are authorized to access it.

In this lab, your task is to perform the following:

Grant the Managers group the **Full Control** permission to the **D:\Personnel** folder.

Remove all inherited permissions that are flowing to the **D:\Personnel** folder.

If a permission appears grayed out, it is an inherited permission. To modify it, you need to disable inheritance and create explicit permissions.

Explanation

Complete this lab as follows:

Open the **Data (D:)** drive.

From the Windows taskbar, select **File Explorer** .

From the left pane, expand and select **This PC** > **Data (D:)** .

Configure NTFS permissions.

From the right pane, right-click **Personnel** and select **Properties** .

Select the **Security** tab.

Select **Edit** .

Select **Add** .

Enter **Managers** as the group that will receive permission to the folder.

Select **OK** .

With the Managers group selected, select the appropriate **Full control** .

Select **OK** .

Prevent inherited permissions from parent.

From the Security tab, select **Advanced** .

Select **Disable inheritance** .

Select **Remove all inherited permissions from this object** .

Select **OK** to close the Advanced Security Settings for Personnel dialog.

Select **OK** to close the Properties dialog.

8.2.7 Practice Questions (Section Quiz)

q_file_sys_sec_dmz_secp8

You have placed a File Transfer Protocol (FTP) server in your DMZ behind your firewall. The FTP server will distribute software updates and demonstration versions of your products. However, users report that they are unable to access the FTP server.

What should you do to enable access?

Answers:

Install a VPN.

Define user accounts for all external visitors.

***Open ports 20 and 21 for outbound connections.**

Move the FTP outside of the firewall.

Explanation:

To allow FTP traffic into your DMZ, you must open the correct ports on the firewall. For FTP, the correct ports are 20 and 21 for outbound connections.

Installing a VPN is not necessary to grant access to external users.

Defining user accounts may be required in some situations, but this scenario requires anonymous access.

Moving the FTP server outside the firewall is not a secure action.

q_file_sys_sec_ftps_secp8

FTPS uses which mechanism to provide security for authentication and data transfer?

Answers:

IPsec

Token devices

Multi-factor authentication

***SSL**

Explanation:

File Transfer Protocol Secure (FTPS) uses Secure Sockets Layer (SSL) to provide security for authentication and data transfer. FTPS is an FTP replacement that brings reasonable security to an otherwise unsecured file transfer mechanism. FTP by itself is unsecured because FTP transmits logon credentials in cleartext and does not encrypt transmitted files.

The following are protocols that are not designed to provide a mechanism to provide secure authentication and data transfer for FTPS:

IPsec - IPsec is a protocol suite for encrypting network communications.

Token devices - A token is a device that employs an encrypted key for which the encryption algorithm - Ã‚Â• the method of generating an encrypted password - Ã‚Â• is known to a network's authentication server.

Multi-factor authentication - Multi-factor authentication (MFA) is a multi-step account login process that requires users to enter more information than just a password.

q_file_sys_sec_ftp_secp8

To transfer files to your company's internal network from home, you use FTP. The administrator has recently implemented a firewall at the network perimeter and disabled as many ports as possible.

Now, you can no longer make the FTP connection. You suspect the firewall is causing the issue.

Which ports need to remain open so you can still transfer the files? (Select two.)

Answers:

443

80

23

***21**

***20**

Explanation:

FTP uses port 21 for connection requests and port 20 for data transfers. Both ports need to remain open for you to transfer files to your company's internal network from home.

Telnet uses port 23, SSL uses port 443, and HTTP uses port 80.

q_file_sys_sec_issuing_assets_secp8

You are the IT manager at a rapidly growing start-up. A new employee has just been hired, and you are responsible for provisioning a user account for the employee. The employee's role involves graphic design, and they will be working remotely.

Which of the following steps is MOST crucial to ensure the employee can effectively perform their duties?

Answers:

Conducting a thorough background check.

Enrolling the employee in a security awareness training program.

***Issuing the employee the necessary hardware and software assets.**

Setting up a VPN for the employee to securely access the company network.

Explanation:

Issuing the employee the necessary hardware and software assets is the most crucial step in this scenario. As a graphic designer, the employee will need specific tools to perform their job. Without these, the employee will be unable to work effectively, regardless of the other provisions in place.

Conducting a thorough background check is an important step in the hiring process and can help ensure the security of the company. However, it is not directly related to the employee's ability to perform her duties as a graphic designer.

Enrolling the employee in a security awareness training program is a good practice and can help protect the company's information. However, without the necessary hardware and software, the employee will not be able to perform their job, regardless of their security awareness.

Setting up a VPN for the employee to securely access the company network is important for security, especially as the employee will be working remotely. However, without the necessary hardware and software, secure access to the company network will not enable the employee to perform their job.

q_file_sys_sec_netbios_secp8

You want to close all ports associated with NetBIOS on your network's firewalls to prevent attacks directed against NetBIOS.

Which ports should you close?

Answers:

67, 68

***135, 137-139**

161, 162

389, 636

Explanation:

NetBIOS uses the following ports:

TCP 135

TCP and UDP 137

TCP and UDP 138

TCP 139

Dynamic Host Configuration Protocol (DHCP) uses ports 67 and 68. Simple Network Management Protocol (SNMP) uses ports 161 and 162. Lightweight Directory Access Protocol (LDAP) uses ports 389 and 636.

q_file_sys_sec_sftp_secp8

A university's IT team must securely transfer large files containing sensitive financial data between two offices in different cities.

Which protocol would be the MOST suitable and secure option for this file transfer?

Answers:

HTTP

FTP

SMTP

***SFTP**

Explanation:

Secure File Transfer Protocol (SFTP) is a secure and encrypted version of File Transfer Protocol (FTP), making it the most suitable and secure option for transferring sensitive financial data between offices. SFTP ensures the confidentiality and integrity of the data during transit.

Hypertext Transfer Protocol (HTTP) is not a secure file transfer protocol, and it transmits data in clear text, which exposes sensitive financial data to potential interception and compromise.

FTP is not secure as it does not use encryption, thus making it unsuitable for transferring sensitive financial data securely.

Simple Mail Transfer Protocol (SMTP) sends and receives emails and is not for secure file transfers.

q_file_sys_sec_shared_folders_secp8

As a senior IT professional, you are tasked with securing your company's file system. The company has a large number of employees who need to access various files and documents for their work.

Which of the following strategies would be the MOST effective in ensuring both security and accessibility of the file system?

Answers:

Implement full-disk encryption on all company computers.

Use a VPN for all data transfers within the company.

***Create shared folders with appropriate access permissions.**

Disable all unused services on company computers.

Explanation:

Creating shared folders with appropriate access permissions is the most effective strategy in this scenario. This allows employees to access the files they need while still maintaining security. Access permissions can be set according to the principle of least privilege, ensuring that employees only have access to the files they need for their work.

Implementing full-disk encryption on all company computers is a good security measure, but it does not directly address the issue of file accessibility for employees. It also may not be necessary for all computers, especially if they do not store sensitive information.

Using a VPN for all data transfers within the company can help secure data in transit, but it does not directly address the issue of file accessibility. Additionally, it might slow down the network and make file access more difficult for employees.

Disabling all unused services on company computers can help reduce potential attack vectors, but it does not directly address the issue of file accessibility. It is a good security practice but does not solve the problem presented in the scenario.

q_file_sys_sec_ssh_secp8

Which of the following file transfer protocols uses SSH to provide confidentiality during the transfer? (Select two.)

Answers:

FTPS

HTTPS

***SCP**

***SFTP**

FTP

Explanation:

Secure Copy Protocol (SCP) and SSH File Transfer Protocol (SFTP) both use SSH to provide confidentiality.

FTPS and HTTPS both use Transport Layer Security (TLS) or Secure Sockets Layer (SSL) to provide confidentiality.

File Transfer Protocol (FTP) is an older TCP/IP protocol that's used for transferring files across systems.

q_file_sys_sec_temporary_privileges_secp8

You are the IT Security Manager at a large corporation. A team of external auditors has been hired to conduct a comprehensive audit of your company's financial systems. They need temporary access to sensitive data and systems.

Which of the following is the MOST appropriate way to handle this situation?

Answers:

Grant the auditors unrestricted access to all systems for the duration of the audit.

***Create new user accounts for the auditors with temporary elevated privileges.**

Grant the auditors the same access level as the company's financial team.

Deny the auditors access to the systems and provide them with printed reports instead.

Explanation:

Creating new user accounts for the auditors with temporary elevated privileges is the most appropriate action. This allows the auditors to access the necessary systems and data while maintaining control over their access levels. Once the audit is complete, these temporary accounts can be deactivated, ensuring no unnecessary access remains.

Granting the auditors unrestricted access to all systems for the duration of the audit is not a good practice. This could expose sensitive data and systems to unnecessary risk. The principle of least privilege should always be applied, even in temporary situations.

Granting the auditors the same access level as the company's financial team may not be appropriate. The auditors may need access to different systems or data than the financial team, and this approach does not take into account the temporary nature of their need for access.

Denying the auditors access to the systems and providing them with printed reports instead is not practical or efficient. It may also hinder the auditors' ability to conduct a thorough and accurate audit.

q_fileperms_acl_secp8

You have a shared folder named Reports. Members of the Managers group have been given write access to the shared folder.

Mark Mangum is a member of the Managers group. He needs access to the files in the Reports folder, but he should not have any access to the Confidential.xls file.

What should you do?

Answers:

Remove Mark Mangum from the Managers group.

***Add Mark Mangum to the ACL for the Confidential.xls file with Deny permissions.**

Configure NTFS permissions for Confidential.xls to allow read-only.

Add Mark Mangum to the ACL for the Reports directory with Deny permissions.

Explanation:

To prevent Mark from accessing one file, edit the ACL for that file, add his user account to the ACL, and configure Deny permissions. The Deny permissions configured on the file override the Write permissions granted to the folder through the group.

Removing Mark from the group would prevent access to the entire folder, not just to the one file.

Configuring deny permissions to the folder for Mark would also prevent access to the entire folder.

q_fileperms_chmod_secp8

You are a system administrator for a Linux server. You have a directory named "project," which contains sensitive data. The directory is currently accessible to everyone in your team.

You need to change the permissions so that only you (the owner) can read, write, and execute, while your team (the group) can only read and execute. Other users should not have any access.

Which of the following commands should you use?

Answers:

chmod 755 project

***chmod 750 project**

chmod 700 project

chmod 777 project

Explanation:

The **chmod 750 project** command is the correct answer. This command gives read, write, and execute permissions to the owner, read and execute permissions to the group, and no permissions to others. This is the desired outcome.

The **chmod 755 project** command is incorrect. This command gives read, write, and execute permissions to the owner and read and execute permissions to the group and others. This is not the desired outcome, as others should not have any access.

The **chmod 700 project** command is incorrect. This command gives read, write, and execute permissions to the owner and no permissions to the group and others. This is not the desired outcome, as the group should have read and execute permissions.

The **chmod 777** project command is incorrect. This command gives read, write, and execute permissions to the owner, the group, and others. This is not the desired outcome, as others should not have any access, and the group should not have write permissions.

q_fileperms_dacl_secp8

You want to give all managers the ability to view and edit a certain file. To do so, you need to edit the discretionary access control list (DACL) associated with the file. You want to be able to easily add and remove managers as their job positions change.

What is the BEST way to accomplish this?

Answers:

***Create a security group for the managers. Add all users as members of the group. Add the group to the file's DACL.**

Add each user account to the file's DACL.

Create a distribution group for the managers. Add all users as members of the group. Add the group to the file's DACL.

Add one manager to the DACL that grants all permissions. Have this user add other managers as required.

Explanation:

Create a security group for the users and add the users to the DACL. A group is an object that identifies a set of users with similar access needs. Microsoft systems have two kinds of groups, which are distribution groups and security groups. Only security groups can be used for controlling access to objects. As manager roles change, add or remove user accounts from the group. Assigning permissions to a group grants those same permissions to all members of the group.

Adding individual user accounts instead of groups to the ACL would require more work as you add or remove managers.

q_fileperms_ntfs_permissions_secp8

A company employee, Jane, is unable to access a shared folder on the company's file server.

The IT department has confirmed that Jane has been granted "Contributor" share permissions for this folder. However, Jane is still unable to access the folder. As a network administrator, you are tasked with troubleshooting this issue.

Which of the following steps would you take first to resolve this problem?

Answers:

***Check if the folder has NTFS permissions set for Jane or a group she belongs to.**

Change Jane's share permissions to "Owner."

Reformat the volume with NTFS.

Assign "Co-owner" share permissions to Everyone.

Explanation:

Checking if the folder has NTFS permissions set for Jane or a group she belongs to is the correct first step. Even if Jane has been granted share permissions, she will not be able to access the folder if no NTFS permissions have been set for her or a group she belongs to. Checking this should be the first step in troubleshooting.

Changing Jane's share permissions to "Owner" would not necessarily resolve the issue if the problem lies with the NTFS permissions. This could potentially give Jane more access than necessary, violating the principle of least privilege.

Reformatting the volume with NTFS is unnecessary and could lead to data loss. The problem is not with the file system itself but with the permissions set within it.

Assigning "Co-owner" share permissions to Everyone is not a good first step. This would potentially give all users full control over the shared folder, which could lead to security risks. It's better to troubleshoot the specific permissions issue first.

q_fileperms_ntfs_secp8

If Mark has a read-write permission to the share \\fileserver\securefiles and a read-only permission to the file coolstuff.docx on the NTFS file system shared by the file share, he is able to perform which action?

Answers:

Delete the file.

***Read the file.**

Rename the file.

Change the contents of the file.

Explanation:

The permissions of the share and file system work together, and the more restrictive of the two is used when accessing the file through the share. In this case, Mark is allowed to read the file.

Because the NTFS permissions are set to read-only, he would not be allowed to delete, rename, or change the file.

q_fileperms_permissions_secp8

You have a file server named Srv3 that holds files used by the development department. You want to allow users to access the files over the network and control access to files accessed through the network or a local logon.

Which solution should you implement?

Answers:

***NTFS permissions and share permissions**

NTFS permissions and file screens

Share permissions and quotas

Share permissions and file screens

Explanation:

Use New Technology File System (NTFS) permissions and share permissions to control access to files. Share permissions apply when files are accessed through the network, and NTFS permissions apply to both network and local access.

Use file screens to restrict the types of files that can be saved within a folder.

8.3 Linux Host Security

As you study this section, answer the following questions:

How do you check for unnecessary network services on a Linux system?

Why is it important to identify open ports? What utility can identify open ports?

Which utility can identify network statistics on a system?

Which commands should you use to disable unneeded daemons?

What are iptables?

In this section, you will learn to:

Remove unnecessary services.

Install and update iptables.

Key terms for this section include the following:

Term	Definition
iptables	iptables is a firewall command line utility for Linux operation systems that uses three policy chains to allow or block network traffic.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.2 Explain common threat vectors and attack surfaces. Removable device 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Patching Hardening techniques Host-based firewall Host-based intrusion prevention system (HIPS) Disabling ports/protocols Removal of unnecessary software
TestOut Security Pro	2.1 Harden physical access 2.1.3 Install and configure a firewall

8.3.1 Linux Host Security (Lesson Video)

Transcript:

The Linux operating system is widely used and preferred in the security community. Just like its Windows counterpart, though, Linux is vulnerable to attacks. By default, the Linux OS is designed to reduce vulnerability and attack surface, but attackers are still able to exploit known vulnerabilities. There's only one way to eliminate all potential threats and that's to disconnect from the internet. Unfortunately, this isn't an option for most users.

With that in mind, you should remain vigilant and take the necessary precautions to protect your system. In this lesson, we'll discuss how to minimize your Linux system's attack surface with simple yet effective methods for reducing the chances of an attack.

The first step is to lessen your system's attack surface as a whole. The more services that are loaded and active on a system, the more opportunities exist for attack. Many services utilize a specific TCP/IP port that's exposed to the network. Services such as DNS, FTP, SNMP, and others may or may not be required on your system. If not, these services should be unloaded and deactivated.

Different Linux distributions load various programs and daemons, and certain programs and daemons are necessary dependencies of others. So make sure you understand your system's requirements before you attempt to unload and deactivate anything. To see a list of all loaded and active services, you can use the â€˜systemctl' command, as shown here. When used, both the service name and description are displayed. Additionally, you can add the â€˜list-dependencies' statement to display a dependency services tree. With this information, you can discover potentially unneeded services and unload them. You first stop the service and then disable it from starting automatically.

A word of cautionâ€”do your research! Don't simply unload a service without knowing what it's used for. It may be a system service or a service that something else depends on. Look up the name in a search engine, or use the â€˜man' or â€˜info' utilities to find out what it's used for. Only then, if you determine it's unneeded, should you stop and disable the service.

It may be necessary to have ports open for a Linux system to work properly. For example, if a server provides DHCP or DNS services or web content, TCP/IP ports must be open to allow everything to work. One way to discover open ports on a system is to use the Nmap utility.

Nmap is a tool that's either bundled with Linux distributions or can be installed using apt or yum. When used appropriately, this tool can provide you with very useful information about a given system or network. When used inappropriately, attackers can take control and perform reconnaissance. For our purposes, we can use Nmap to show us our system's open ports.

Often, open ports are indications that a service is running on the system that controls that port. If we see a port open that shouldn't be, it can tell us that a service is running on the system that we probably don't need. We can check the system and stop and disable the service if needed.

In the example here, Nmap shows the default gateway's open ports. There may be ports we expect and others we don't. For example, here we see that port 80 is open with the service http. If this machine doesn't require web services, we may want to close that port by stopping the web server. If we see an unexpected open port, we need to find the service that's opening that port to determine whether the service is necessary or not. If it isn't, we can stop and disable the service to close the port. Just use this tool with caution. Since Nmap is widely used by hackers, running it may trigger alarms and alerts.

There are many command line parameters available with Nmap that let us customize the desired command output. Let's discuss a few. The â€˜â€“sU' and â€˜â€“sT' parameters display only UDP and TCP ports. Keep in mind that these parameters might make your scan take an extraordinarily long time. Use â€˜-6' to choose IPv6 scanning. To enable OS detection, use â€˜-A' or
â€˜â€“O'. We use the â€˜-p' parameter to specify a range, while the â€˜-sn' parameter enables a ping-only scan to determine if a host is up without having to perform a port scan.

Netstat is another Linux or Windows tool that we can use to scan for open ports. In this example, a Linux machine was used to display open ports on a system with IPv4. I used the â€˜â€“l' option to view a list of listening sockets and the â€˜4' option to show only IPv4 addresses on the Linux host.

Here you see I have a socket open for ssh. Also, the system is running NTP. These are TCP/IP ports that are well known as ports 22 and 123, respectively. By using this output, I can also see what else is running on the system.

We can customize the Netstat utility to provide us with only the information we're specifically looking for. As with Nmap, the Netstat utility has several command line parameters. For example:

â€˜-a' lists all listening and non-listening sockets.
â€˜-I' displays statistics for all your network interfaces.
â€˜-l' displays just listening sockets or sockets that are open for listening.
â€˜-s' displays summary information for each protocol enabled on the system.
â€˜-r' shows your routing table.

Keep in mind that these options may vary in availability depending on whether we're using the Linux version of Netstat or the Windows version.

In addition to unloading unneeded services, you also need to make sure that your Linux operating system remains current by installing operating system and program updates. For example, as software is released, it's usually not perfect and free of defects. Defects can take the form of logic errors in which a program doesn't calculate or apply input correctly. There could be something wrong with a screen display as well. In addition, vulnerabilities can be introduced

that would allow threat actors to compromise your system. Often, these defects are found only after software has been released. So once the software is in use and the defects are found and corrected, it's important that you update your software immediately.

You can manually update the packages on your system using package managers such as yum or apt. The one you use will depend on which distribution you're using. These processes scour the repository and search for newer updates and patches. If a single package requires an update, the package name can be specified to check if that package requires an update. Most often, though, it's a good practice to update all available package updates. For example, if you're running a Red Hat variant, such as CentOS or Fedora, run the â€˜yum update' command, as shown here. In this example, you see that only a single packageâ€”libexifâ€”needs an update.

When you use a Debian variant like Ubuntu or Mint, you first run the â€˜apt update' command. This refreshes the repository information. After that, you run â€˜apt upgrade' to perform the manual update process. Notice that in this case, there are 28 packages to be upgraded. An exception to utilizing apt or yum is with SuSE Linux. SuSE Linux uses a different package manager called zypper. The syntax for the command is similar to apt and yum and accomplishes the same tasks.

The last topic for enabling Linux host security is to ensure a host-based firewall is running on your Linux system. One philosophy of computer security is to utilize a layered approach. Most systems are protected first by a network firewall. If anything does happen to get through, threat actors also have to contend with the host system's firewall as well. Most, if not all, Linux systems are distributed with host-firewall software. This software is known as firewalld.

In some instances, there might even be a series of firewalls that are specific to protecting certain network aspects. These together form a solid barrier to attackers. But nothing is foolproof and adding another layer of protection is always a good idea. The Linux host-based firewall adds additional protection from outside entities, and it can also protect the system from insider threats, too.

A host-based firewall acts like a gatekeeper between your system and the external and internal network. This firewall monitors all traffic that flows in both directions between the computer and network. You configure this firewall with a list of rules called access control lists, or ACLs. These rules define what is and isn't allowed to pass through. Several Linux distributions use the firewalld daemon to implement host-based firewalls. Other distributions might use a different package, so check your distribution vendor to see which one to use.

First, make sure firewalld has been installed on your system. You can use the â€˜yum' or â€˜apt' command lines to check whether it has, as shown here. If those commands return a value, you know firewalld has been installed and you can begin configuring it. If you get a blank message or one stating firewalld could not be found, you'll need to download and install it. This can also be done using the â€˜yum' or â€˜apt' commands. Know that firewalld depends on the Python programming language. If not already installed, just note that the installer will also install the needed Python dependencies on your system.

If firewalld is already installed or if you just installed it, you need to ensure it's running and active with the â€˜systemctl' command shown here. It might not be running, and you might see inactive or dead as the status. If that's the case, you can start the daemon by issuing the â€˜systemctl start firewalld' command. Checking the status again should show the daemon running.

With the firewall running, there are several different commands that you can use to manage it.

We use the â€˜firewall-cmd' command and add parameters to control how the firewall operates and is configured. A sampling of the commands are listed here. Run the â€˜firewall-cmd --state' command to check the firewall's status. You can also run â€˜firewall-cmd --get-active-zones' to display the default firewall zone configuration.

This particular firewall package comes with several predefined firewall zones that you can use, starting at extremely secure to not very secure at all. Predefined zones include Home, External, Work, and Trusted. The default zone most distributions predefine is Public, but you can actually pick whichever one you want to use.

Once you have the firewall running and a default zone set up, there may be situations in which you need to open a particular firewall port to support a particular service. For example, if the system is set up as a web server, you must let web traffic through. If you don't configure firewall exceptions, the web traffic will be blocked.

Web servers typically use ports 80 and 443 or HTTP and HTTPS for web traffic. We need to allow these ports to remain open so that the web server is reachable. To do this, we need to allow the right ports and protocols through the firewall. To enable web services, we open ports on the firewall. To open ports, we add them to the exception list on the firewall itself.

This is also done using the â€˜firewall-cmd' command. For example, we first use the â€˜--permanent' parameter to store the exception in the database so that it remains persistent even if the computer is restarted. Next, we need to define which zone we're using and then the port we want open.

It's also possible to use the protocol rather than the port number when placing exceptions into the firewall. We place the exception in the firewalld configuration with HTTPS by using the same command but substituting the protocol. The last thing we need to do is restart the firewall or reload the configuration. If the system is in use and we don't want to disrupt

it, we enter the â€˜reload' command, which allows the firewall daemon to continue running. The new configuration is loaded and activated for us in the background.
The last thing we need to do is restart or reload the newly configured firewall. This allows us to use these changes immediately without rebooting the system. The â€˜reload' parameter leaves the firewall activated.
That's it for this lesson. In this lesson, we discussed some procedures to keep a Linux host safer from attacks. We discussed removing unneeded services from your system and using Nmap and Netstat to identify running services. We also discussed how keeping the system current with OS and program updates is important. Finally, we ended by discussing the benefits of implementing the host-based firewall called firewalld.

8.3.2 Removing Unnecessary Services (Demo Video)

Transcript:

It's important to know many things about services: how to start them, see if they're running, stop them, restart them, enable them, and disable them. For this demo, we're going to work with an Ubuntu Linux system and practice working with some services.

I'm already logged on to the system, and I'll start the terminal. I have a shortcut here, but you can also go here, to Show Applications, and then search for 'terminal' in the search box.
I want to see the init.d directory. We're interested in services that are running. Let's go to the directory by typing 'cd /etc/init.d/'. Press Enter. Now I'll do an 'ls' to see what's in this directory.
All right, here are my services. Here, I have Apache server. Here, I have mysql. Over here, we have ssh. And down here, there's FTP. On this system, there's no reason to have FTP running, so let's see how we can stop it.
I'll type 'cd' to get back to root and 'clear' to clean up my screen.
So far, I've used init.d to view the services. To manage the services, I'm going to use the systemctl command. Using init.d scripts to manage services isn't possible on newer systems anymore, but I still use it to view services sometimes.
Let's look at my FTP server. First, I want to see its status. So, let's type 'sudo systemctl status vsftpd' and press Enter. It says it's enabled, but it's currently stopped.
I'm going to use my up arrow to get to the last command, change it from status to 'start', and press Enter. Now I'll arrow up until I see status and press Enter again.
I can see that it's running along with some statistics, such as how long it's been up. I can also see the process ID, or PID, right here. Let's clear the screen.
Now I'm going to use my up arrow to go back to my last command. I'm going to backspace here, get rid of the word "start:", and type in 'stop' to stop the FTP service.
I'll arrow up again until I get to status and press Enter. Now you can see that the process has stopped. I'll press Ctrl+C to exit this command and then type 'clear'.
Now I might say, oops, that's not what I wanted to do! It's okay. We'll just arrow up again until we get to start. Press Enter, and everything's just fine again.
What if I made changes to my FTP server, and I need to restart the service? I can just do a restart. Let's up arrow to the last command, backspace here, type in 'restart', and press Enter. Arrow up to the last command, backspace, and change this to 'status'. Press Enter, and you can see my service restarted. I'll clear the screen again.
So far, we've stopped and started our FTP service. But what if I want to have it start at system startup? Or, perhaps, not start when the system starts up?
First, let's see if the service is enabled or not. I can do that with the is-enabled command. I'll arrow up to the last command and change this part to 'is-enabled'. Press Enter, and right here, it says that it's enabled.
What if you want to have your service start every time your system starts? You can do that by changing this to 'enable' and pressing Enter. Now you can see the status is enabled.
If I don't want the service to start, I'll arrow up and change this to 'disable'. Press Enter, and that'll tell this system not to start FTP at startup. I'll arrow up and change this to 'status', and you can see that, up here, it says "disabled", so when the system boots, FTP won't start. Note that the service is still running at the moment because all we did was keep it from running at startup; we didn't stop it from running at this moment.
That's it for this demo. In this demo, we worked with system services. First, we viewed the services to see what was on our system. Then we practiced stopping, starting, enabling, and disabling services.

8.3.3 Linux Host Security Facts

This lesson covers the following topics:

- Protecting ports
- Host-based firewalls and IPS
- Hardening a Linux system

Protecting Ports

Physical device port hardening involves restricting the physical interfaces on a device that can be used to connect to it, thereby reducing potential avenues of physical attack. One common technique is disabling unnecessary physical ports such as USB, HDMI, or serial ports when they serve no business purpose. Doing so helps prevent unauthorized data transfer, installation of malicious software, or direct access to a system.

Port control software provides additional protection by only allowing authorized devices to connect via physical ports based on device identifiers. For instance, it might block all USB mass storage devices except company-approved ones.

Security analysts can leverage settings in device firmware or UEFI/BIOS for port hardening to disable physical ports or to require a password before a device can boot from a nonstandard source like a USB drive. For devices such as tablets and laptops that depend upon wireless protocols, disabling the automatic network connection feature can prevent the device from using potentially insecure or rogue networks.

As revealed by researcher Karsten Nohl in his BadUSB paper (https://assets.website-files.com/6098eeb4f4b0288367fbb639/62bc77c194c4e0fe8fc5e4b5_SRLabs-BadUSB-BlackHat-v1.pdf), exploiting the firmware of external storage devices, such as USB flash drives and even standard-looking device charging cables, presents adversaries with an incredible toolkit. The firmware can be reprogrammed to make the device look like another device class, such as a keyboard. In this case , it could then be used to inject a series of keystrokes upon an attachment or work as a keylogger. The device could also be programmed to act like a network device and corrupt name resolution, redirecting the user to malicious websites.

Another example is the O.MG cable (theverge.com/2019/8/15/20807854/apple-mac-lightning-cable-hack-mike-grover-mg-omg-cables-defcon-cybersecurity), which packs enough processing capability into an ordinary-looking USB-Lightning cable to run an access point and keylogger.

A modified device may have visual clues that distinguish it from a mass-manufactured thumb drive or cable, but these may be difficult to spot. You should warn users of the risks and repeat the advice to never attach devices of unknown provenance to their computers and smartphones. If you suspect a device is an attack vector, observe a sandboxed lab system (sometimes referred to as a sheep dip) closely when attaching the device. Look for command prompt windows or processes, such as the command interpreter starting and changes to the registry or other system files.

Not all attacks have to be so esoteric. USB sticks infected with ordinary malware are still incredibly prolific infection vectors. Hosts should always be configured to prevent autorun when USB devices are attached. USB ports can be blocked altogether using most types of host i ntrusion d etection s ystems (HIDS).

Protect logical ports by implementing measures to secure and control access to ports within a computer system or network. Logical ports are software-based communication features that enable data exchange between applications or services. Common examples of logical ports include the well-known ports used by TCP/IP and UDP protocols.

Firewalls protect logical ports by examining network traffic and enforcing security policies to allow or block specific connections based on port numbers, source and destination addresses, and protocols. Service hardening practices ensure that services running on logical ports are hardened against security threats. Examples include keeping software updated and turning off unnecessary services.

Host-Based Firewalls and IPS

Host-based firewalls and intrusion prevention systems (IPS) are vital elements of endpoint hardening, as they provide controls for incoming and outgoing network traffic and are essential for detecting potential attacks. An important technique for using them when hardening endpoints involves implementing default-deny policies to block all traffic unless explicitly allowed. This tactic ensures that only approved services and applications can communicate. Configuring firewalls to block or allow traffic based on port numbers is also critical to minimize entry points for attack. Traffic filtering enables firewalls and IPS to sift through traffic based on parameters like IP addresses, protocols, and services to block malicious traffic or only allow traffic to use secure protocols.

An integral part of IPS is detecting and preventing intrusions by monitoring for known malicious patterns or anomalies in network traffic. Advanced host-based firewalls often include application control features that permit only trusted applications to communicate. The logs generated by host-based firewalls and IPS support rapid detection and response when integrated with other security tools like security information and event management (SIEM) systems.

Hardening a Linux System

The following table describes the general procedures for increasing endpoint security of a Linux system:

Security Task	Procedure
Remove unnecessary software	Unnecessary software occupies disk space and could introduce security flaws. To remove unnecessary software: Enter one of the following commands: **yum list installed** or **dnf list installed** to see installed RPM packages on the computer. **apt** **apt autoremove** automatically removes unused packages **apt list** list all installed packages **dpkg get-selections** to see installed Debian packages on the computer. Research the function of any unrecognized package to determine if it is necessary. Use one of the following commands to uninstall unnecessary packages. **yum erase *packagename*** **dnf remove *packagename*** **apt remove *packagename***

	rpm -e *packagename* **dpkg -r *packagename***
Check for unnecessary network services	Unnecessary network services waste computer resources and increase the system's attack service. To remove unnecessary network services: Find all installed services and determine which are not needed: DNS, SNMP, DHCP, and others. **systemctl --type=service --state=active** Use the **man** command and the internet to research services you do not recognize. If the service is not needed, determine if it is a dependency for another service. Disable the service by using the following command: **systemctl disable *servicename*** Use one of the following commands to stop the script immediately: **systemctl stop *servicename*** Use one of the following commands to remove the script package entirely: **yum erase *packagename*** **dnf remove *packagename*** **apt remove *packagename*** **rpm -e *packagename*** **dpkg -r *packagename***
Locate open ports	Open ports can provide information about which operating system a computer uses. Also, they can provide entry points or information about ways to formulate an attack. To locate open ports: Install the **nmap** utility if it is not already installed. **yum install nmap** **dnf install nmap** **apt -i nmap** Use both of the following commands to scan for open ports:

	nmap -sT *ipaddress|fqdn* scans for TCP ports **nmap -sU** *ipaddress|fqdn* scans for UDP ports Determine which services use the open ports. Disable any unused service using the open ports information. (*Make sure the service used is not a dependency for another service* .) **systemctl disable *servicename*** **systemctl stop *servicename***
Check network connections	Open network connections (open sockets) on a computer create a security risk. A *socket* is an endpoint of a bi-directional communication flow across a computer network. Use the following **netstat** (network statistics) or **ss** (socket statistics) options to identify the open network connections on Linux systems: **-a** lists both listening and non-listening sockets. **-l** (*lowercase 'L'*) lists listening sockets. **-s** displays statistics for each protocol. **-i** displays a table of all network interfaces.

8.3.4 Configure iptables (Demo Video)

Transcript:

In this demo, we're going to look at iptables, the default firewall for Linux systems. You configure iptables from the from the command line. We aren't going to cover everything about iptables because that's a vast subject that goes way beyond the scope of this course, but we are going to introduce the topic.

We know that a firewall is a part of a computer system, or network, that is designed to block unauthorized access while permitting authorized communications to hosts. Let's start by talking about the filtering table. A filter table has three of what's know as chains. The input chain applies rules to packets coming into the system. The forward chain is for packets being routed through the system. And the output chain is for locally generated packets sent out from the system.
Let's look at our iptables. I have a Linux system running here, and I want to open the terminal, or command line. I'll select it. My terminal application is right here. But if it wasn't listed, I could type 'terminal' in the Search box.
iptables is almost always preinstalled on Linux systems. But if it is not, we can run a command to install or update it. If it is installed, running this command won't hurt anything, so let's go ahead and do that now. To install or update iptables, we have to run this as a root user, so we need to use the sudo command. So, let's type in 'sudo apt-get install iptables'. Linux prompts me for my password, so I'll type that in right here. We can see it tells us that iptables is already installed, and I have the latest version. We are good to go here.
The first thing I want to do is see what's currently configured. Remember that Linux is case sensitive, so pay attention to that when you type in commands. I'm going to clear the screen by typing .clear.. Let's type in 'sudo iptables -L'. Notice that the policy is set to Accept for all three chains.
Let's say our Linux system is getting a very large number of ping requests, so we decide to block echo requests. We can do this with iptables. Let's jump over to a Windows 10 workstation. I'm going to ping my Linux system just to be sure they

can communicate. I know that the IP address for the Linux system is 192.168.0.7, so let's type 'ping 192.168.0.7' and press Enter. You can see I get a response. Now let's jump over to our Linux system and block pings.
Okay, I'm back on the Linux system. I'm going to clear the screen again. To block pings, I need to enter the following command: 'sudo iptables -A INPUT -p icmp --icmp-type echo-request -j DROP'. Remember that this is case sensitive. With this command, the -A tells it I want to add a rule, the -p specifies the protocol. There are a whole lot of switches you can use, but we only use a couple here. We also told iptables to drop the request. We also could have told it to reject the request. The difference is that if we drop it, the system responds as if there isn't even a system there. It does nothing except drop the request. Reject, on the other hand, responds with a message saying the destination port is unreachable. That might tell a hacker that there is something there, but it's not replying.
Now let's go back to our Windows 10 workstation once more and try to ping again. I'm going to type .ping 192.168.0.7. and press Enter. You see that I get request timed out messages back, so the rule is working.
That's it for this demo. In this demo, we talked about iptables. We ran the command to install and update iptables. We looked at the iptables default policies and then configured a simple rule to drop ping requests to our Linux system.

8.3.5 Configure iptables Facts

This lesson covers the following topics:

- Chains
- Actions performed
- Example iptables commands

Chains

The Linux iptables firewall utility uses *policy chains* (sets of rules) to allow or block network traffic. When a connection is initiated to your system, iptables looks for a matching rule. If it doesn't find one, it uses the default action in the tables. Be aware that iptables almost always come pre-installed on any Linux distribution.

The filter table in iptables has three chains. The following table describes them.

Chain	Description
Input	This chain controls the behavior of incoming connections. For example, if a user attempts to ping the system, iptables attempts to match the IP address and port to a rule in the input chain.
Forward	This chain is used for packets leaving the system. These are incoming connections that aren't delivered locally. In other words, the traffic is not destined for the router; the router forwards the traffic to the destination device.
Output	This chain is used for outgoing connections. For example, if you ping testout.com, iptables check its output chain to see what the rules are regarding ping and testout.com before allowing or denying the ping request.

Actions Performed

You can accept, drop, or reject the connections. After you define your accept rules, you should create a rule to drop all other traffic to prevent unauthorized access to the system.

Action	Result
Accept	Allows the connection.
Drop	Drops the connection. For example, an IP address in a rule with a drop action pings your system; the request is dropped. No response is sent to the user.
Reject	Rejects the connection but will send a response back. This lets the sender know that the traffic reached a system but was rejected.

Examples iptables Commands

The following table describes commands for iptables. Keep in mind that these are only a few examples; there are many more iptable commands.

Action	Command
List current rules	**sudo iptables -L**
Clear current rules	**sudo iptables -F**
Save iptables changes (Ubuntu)	**sudo /sbin/iptables-save** The command may be different on other Linux systems.
Drop all incoming traffic	**sudo iptables -A INPUT -j DROP**
Block connections from 192.168.0.254	**sudo iptables -A INPUT -s 192.168.0.254 -j DROP**
Block SMTP mail on port 25	**sudo iptables -A OUTPUT -p tcp --dport 25 -j REJECT**
Allow SMTP mail on port 25	**sudo iptables -A INPUT -p tcp --dport 25 -m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT** **sudo iptables -A OUTPUT -p tcp --sport 25 -m conntrack --ctstate ESTABLISHED -j ACCEPT**
Allow HTTP traffic on port 80	**sudo iptables -A INPUT -p tcp --dport 80 -m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT** **sudo iptables -A OUTPUT -p tcp --sport 80 -m conntrack --ctstate ESTABLISHED -j ACCEPT**

	To allow HTTPS, you would use port 443.
Allow HTTP traffic on port Allow HTTPS traffic on port 443	**sudo iptables -A INPUT -p tcp -m multiport --dports 80,443 -m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT** **sudo iptables -A OUTPUT -p tcp -m multiport --dports 80,443 -m conntrack --ctstate ESTABLISHED -j ACCEPT**

8.3.6 Practice Questions (Section Quiz)

q_linux_host_sec_firewalls_secp8

As a network administrator, you are tasked with implementing measures to secure and control access to logical ports within your organization's computer system. You are considering several strategies to accomplish this task.

Which of the following would be the MOST effective method to protect logical ports?

Answers:

Implementing service hardening practices

Using secure protocols for data exchange

***Implementing firewalls**

Regularly updating software

Explanation:

Implementing firewalls is the most effective method for protecting logical ports. They examine network traffic and enforce security policies to allow or block specific connections based on port numbers, source and destination addresses, and protocols. This directly protects the logical ports from unauthorized access and attacks.

While service hardening is an important part of overall security, it primarily focuses on ensuring that services running on logical ports are hardened against security threats. It does not directly protect the logical ports themselves.

Using secure protocols is a good practice to ensure the security of data in transit. However, it does not directly protect the logical ports from unauthorized access or attacks.

Regular software updates are important to fix security vulnerabilities. However, they do not directly protect the logical ports from unauthorized access or attacks.

q_linux_host_sec_ips_solution_secp8

As a cybersecurity analyst, you are tasked with enhancing the endpoint security of your organization's network. You are considering several strategies to accomplish this task.

Which of the following would be the MOST effective method to control incoming and outgoing network traffic and detect potential attacks?

Answers:

***Implementing host-based firewalls and IPS with default-deny policies**

Regularly updating all software and applications

Implementing a strong password policy

Regularly conducting penetration testing

Explanation:

Implementing host-based firewalls and IPS with default-deny policies is the most effective method, as it provides controls for incoming and outgoing network traffic and is essential for detecting potential attacks. Default-deny policies block all traffic unless explicitly allowed, ensuring that only approved services and applications can communicate.

While regularly updating all software and applications is a good practice to fix security vulnerabilities, it does not directly control network traffic or detect potential attacks.

While a strong password policy is important for securing user accounts, it does not directly control network traffic or detect potential attacks.

While penetration testing is important for identifying vulnerabilities in the network, it does not provide real-time control of network traffic or detection of potential attacks.

q_linux_host_sec_netstat_secp8

Which command should you use to display listening and non-listening sockets on your Linux system? (Tip: enter the command as if in Command Prompt.)

Answers:

netstat -a

netstat -a

Explanation:

Use **netstat -a** to identify listening and non-listening sockets on a Linux system. A socket is an endpoint of a bidirectional communication flow across a computer network. Be aware of the other common **netstat** options:

-l lists listening sockets.

-s displays statistics for each protocol.

-i displays a table of all network interfaces.

q_linux_host_sec_nmap_secp8

Which command should you use to scan for open TCP ports on your Linux system? (Tip: enter the command as if in Command Prompt.)

Answers:

nmap -sT

nmap -sT

Explanation:

Use **nmap -sT** to scan for open TCP ports. Open ports can provide information about which operating system a computer uses and might provide entry points or information about ways to formulate an attack.

Use **nmap -sU** to scan for open UDP ports.

q_linux_host_sec_open_secp8

You need to increase the security of your Linux system by finding and closing open ports.

Which of the following commands should you use to locate open ports?

Answers:

netstat

***nmap**

traceroute

nslookup

Explanation:

Use **nmap** to locate open ports. Open ports can provide information about which operating system a computer uses and might provide entry points or information about ways to formulate an attack. Use one of the following commands to scan for open ports:

nmap -sT scans for TCP ports.

nmap -sU scan for UDP ports.

The **netstat** command shows the status of listening and non-listening sockets. A socket is an endpoint of a bidirectional communication flow across a computer network.

The **nslookup** command is used for name resolution requests.

The **traceroute** command tests and displays connectivity between devices.

q_linux_host_sec_service_secp8

What does the **netstat -a** command show?

Answers:

All listening sockets

All connected hosts

***All listening and non-listening sockets**

All network users

Explanation:

The **netstat -a** command shows the status of all listening and non-listening sockets.

q_linux_host_sec_uefi_bios_secp8

As a security analyst at a large corporation, you have been tasked with implementing physical device port hardening techniques to reduce potential avenues of physical attack. You are considering several techniques to accomplish this task.

Which of the following would be the most effective method to prevent unauthorized data transfer, installation of malicious software, or direct access to a system?

Answers:

Disabling unnecessary physical ports such as USB, HDMI, or serial ports

Implementing port control software

***Leveraging settings in device firmware or UEFI/BIOS**

Disabling the automatic network connection feature

Explanation:

Leveraging settings in device firmware or UEFI/BIOS is the most effective method. It allows you to disable physical ports or require a password before a device can boot from a nonstandard source like a USB drive. This can prevent unauthorized data transfer, installation of malicious software, or direct access to a system.

Disabling unnecessary physical ports such as USB, HDMI, or serial ports: This is a good practice as it reduces the number of potential entry points for an attacker. However, it does not prevent an attacker from using the remaining enabled ports to launch an attack.

Implementing port control software: This can provide additional protection by only allowing authorized devices to connect via physical ports based on device identifiers. However, it does not prevent an attacker from spoofing the identifiers of an authorized device.

Disabling the automatic network connection feature: This can prevent a device from using potentially insecure or rogue networks. However, it does not prevent an attacker from physically connecting to the device and launching an attack.

q_linux_host_sec_yum_secp8

You want to make sure no unneeded software packages are running on your Linux server.

Select the command from the drop-down list that you can use to see all installed RPM packages.

Answers:

yum list installed

yum list packages

yum list rpm packages

yum list installed

yum list rpm installed

Explanation:

Unneeded software takes disk space and could introduce security flaws. To see all the RPM packages installed on your Linux server, run the following command:

yum list installed

After running this command, complete the following:

Research the function of any unrecognized RPM package to determine whether it is necessary.

Use **yum** or **rpm** to uninstall unneeded packages.

q_conf_iptables_drop_secp8

Which action would you use in a rule to disallow a connection silently?

Answers:

***Drop**

Accept

Forward

Reject

Explanation:

The Drop action is used to disallow a connection silently; the sending system receives no notice. The Reject action also disallows a connection but sends a TCP RST packet or an ICMP port unreachable packet back to the system that sent the original packet.

Accept would allow the packet.

Forward is a chain, not an action in iptables.

Reject rejects the connection but will send a response back.

q_conf_iptables_input_secp8

In which of the iptables default chains would you configure a rule to allow an external device to access the HTTPS port on the Linux server?

Answers:

Forward

Output

***Input**

Accept

Explanation:

The Input chain would be where you would place the rule as it is used for inbound connections.

The Output chain is for outbound connections.

The Forward chain sends connections through the Linux server to another device.

The Accept action can be used in a rule to allow a connection. However, it is not a chain.

q_conf_iptables_rst_secp8

Which type of packet would the sender receive if they sent a connection request to TCP port 25 on a server with the following command applied?

sudo iptables -A OUTPUT -p tcp --dport 25 -j REJECT

Answers:

ICMP Unreachable Port

ACK

SYN

***RST**

Explanation:

Because the packet is TCP and is blocked by the Reject action, the server would send a TCP RST packet back to the sender.

ICMP Unreachable Port is sent by iptables if a UDP packet is blocked by the Reject action.

A SYN packet would indicate that the server is proceeding with the connection, which would not happen with the Reject action. If it were allowed, the ACK would generally be sent with the SYN to acknowledge the initial connection while the SYN starts the next part of the TCP three-way handshake.

q_conf_iptables_smtp_secp8

You have configured the following rules. What is the effect?

sudo iptables -A INPUT -p tcp --dport 25 -m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT
sudo iptables -A OUTPUT -p tcp --sport 25 -m conntrack --ctstate ESTABLISHED -j ACCEPT

Answers:

Block SSH traffic

Allow SSH traffic

***Allow SMTP traffic**

Block SMTP traffic

Explanation:

These rules would allow inbound and outbound Simple Mail Transfer Protocol (SMTP) connections on TCP port 25, the default port for SMTP.

These rules use the Accept action, so they would not block SMTP or Secure Shell (SSH).

SSH is on TCP port 22, so these rules would not affect SSH.

q_conf_iptables_sudo_secp8

Which command would you use to list all of the currently defined iptables rules?

Answers:

sudo iptables -F

sudo iptables -A INPUT -j DROP

sudo /sbin/iptables-save

***sudo iptables -L**

Explanation:

sudo iptables -L lists all of the currently defined rules.

sudo iptables -A INPUT -j DROP would drop all incoming traffic.

sudo /sbin/iptables-save saves changes to iptables on Ubuntu.

sudo iptables -F would flush all current rules from iptables.

8.4 Wireless Overview

As you study this section, answer the following questions:

Which device broadcasts information and data over radio waves?

What are the two modes of wireless network configuration?

Where is a Wireless LAN Controller (WLC) installed?

In this section, you will learn to:

Configure a wireless connection.

The key terms for this section include:

Term	Definition
Service set identifier (SSID)	A unique name that identifies a wireless network.
Wireless access point (WAP)	A wireless access point broadcasts information and data over radio waves.
Wireless interface	The interface in a device, such as a laptop or smart phone, that connects to the wireless access point.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.1 Given a scenario, apply common security techniques to computing resources. Wireless devices Installation considerations Site surveys Heat maps Wireless security settings Wi-Fi Protected Access 3 (WPA3) AAA/Remote Authentication Dial-In User Service (RADIUS)
TestOut Security Pro	2.2 Harden Network Devices

	2.2.2 Configure and Access a Wireless Network 2.2.4 Harden a Wireless Network

8.4.1 Wireless Networking Overview (Lesson Video)

Transcript:

Wireless networking is now commonplace in most network environments. In fact, a lot of networks use it as the primary connection method. Chances are you probably have a wireless network in your home. In fact, setting up a residential wireless network at home is pretty simple. Even setting up the SSID broadcast and encryption protocol on the access point is very straightforward.

In an enterprise environment, setting up a wireless network isn't as simple. This is because there are additional considerations that need to be taken into account in an enterprise environment. But before we discuss the impact of wireless devices in an enterprise network, let's highlight a couple of things with consumer-grade home routers.
You probably only have a single wireless access point at home. This device is actually three devices in one: a router, a switch, and a wireless access point.
The router connects your home local area network to the ISP's network. The switch allows you to connect several network devices together using an Ethernet cable, most likely unshielded twisted pair. The wireless access point allows for wireless connections to the local area network.
An enterprise wireless network is different from a home wireless network in at least three different ways. First, instead of having a single all-in-one device, equipment is typically separated into distinct components that perform a single function. For example, a router only performs routing operations, a switch only performs switching operations, and a wireless access point only acts as a bridge between the wireless and wired network.
Second, enterprises need to support many more devices and cover a much larger area than home networks. Some enterprises need to support thousands of clients and cover entire business complexes; for example, many universities have wireless connectivity on the entire campus.
And third, devices need to be able to move around the enterprise without any connection troubles. The transition between different access points needs to be seamless.
In an enterprise network, this is done by using multiple access points, which are scattered throughout the campus to provide adequate coverage. The access points will typically all connect to the same local area network through a switch. Now, each of these access points can only support a limited number of clients, so having multiple access points also allows each access point to share the client load. For example, if this switch has the maximum number of connections, new wireless devices that require network access will be connected to this access point, here, instead. This is possible because the range of the access points here overlap, which actually brings up another important point: overlapping signals can cause interference.
It's also possible for devices to constantly bounce between connecting to this AP and this AP, and so on. This could be really frustrating. In addition, if any of these configuration settings on either of these APs is different, then this can cause connection issues, and the more APs you add to a network, the more devices you need to manage.
As you can see, a wireless network can quickly become a management nightmare. Luckily, enterprise networks use a device called a wireless LAN controller (WLC) that manages all the access points in a network. The typical enterprise wireless controller will be in a networking closet and connected to a switch. The controller can then communicate with and manage each of these APs.
In addition to managing configuration settings, the controller can also manage client connections and manage AP loads. In essence, a wireless controller allows multiple APs to work together as a single system, instead of each working in isolation.
It's important to know that a wireless controller by itself doesn't actually produce any type of wireless signal--that's the job of the AP. The controller only manages the APs.
Another type of wireless device is known as a light-weight access point, or LWAP.
LWAPs are used in conjunction with a wireless controller. LWAPs contain very little embedded intelligence or technology. They are connected to the wired network via a network cable and are able to communicate directly with the WLC, which handles things like client connections, authentication, propagating configurations, et cetera.
Let's next look at the two ways a wireless network can be configured. The first one is ad hoc mode.

In ad hoc mode, also called peer-to-peer mode or wireless mesh mode, there is no WAP or WLC. Instead, two or more wireless devices connect directly with each other to send and receive information. The more devices that connect, the more individual connections that are made. And you can see where the term wireless mesh comes from when you see all these individual connections.
The second mode is called infrastructure mode. In infrastructure mode, a workstation communicates with a notebook by sending the signal to the WAP which transfers it to the notebook system. The infrastructure mode is more scalable than ad hoc mode. The infrastructure mode also makes it easy to connect the wireless network to a wired network. It is not easy to do that with an ad hoc implementation.
Once you establish the wireless connection between the network interface card and the WAP, the wireless network operates similar to traditional wired Ethernet networks.
Now, it's important to know that wireless networks use what's called an unbounded medium to communicate. That is, they send signals in all directions that can be picked up by any device in range they aren't bounded by insulated wires like wired networks. Because of this, wireless networks are inherently less secure than a bounded, wired network, so you need to take extra precautions when securing a wireless network.

8.4.2 Wireless Installation (Lesson Video)

Transcript:

When you set up a wireless network, you can't just put up a bunch of wireless access points and call it a day. There are different installation considerations you'll need to look at before installing your devices. The first step is to understand and define the usage needs of your wireless network. Once you know them, it's time to conduct a site survey. In this lesson, we're going to discuss the site survey process.

A site survey ensures that the wireless network will perform as desired. It examines the physical layout and barriers of the location, which lets you determine where to install the access points. Multiple site surveys will be completed during the installation process.
There are three different types of site surveys performed at different times in the installation process: passive surveys, active surveys, and predictive surveys. If a survey isn't conducted properly, your network won't perform, and that will cost a lot more time and money than performing surveys correctly.
To conduct a survey, you'll need to use a Wi-Fi analyzer. A dedicated Wi-Fi analyzer tool can provide an in-depth look at the wireless signals in the area being tested. There are also several Wi-Fi analyzers programs you can run from a laptop, tablet, or smartphone.
Before you perform a survey, you need a diagram or map of your location. It's important to make sure that the map is properly scaled and labeled so you can get a proper reading. Many analyzer tools allow you to import the map so you can overlay the survey results.
The initial survey is a passive survey. This survey is performed without the analyzer connecting to any specific wireless access point. It's just in a listen-only mode. Temporary WAPs are placed, and the surveyor walks around with the analyzer to measure signal strength and interference levels.
Once the passive survey has been completed, a heat map is generated. A heat map shows Wi-Fi signal strength. They usually show strong signals in green, medium levels in yellow, and weak spots in red.
You may need to perform multiple passive surveys to determine the best placement for your access points. You also need to make sure that the WAPs are installed in secure locations so people can't physically access them. Once you've figured out where to install the WAPs and installed them, you should conduct an active survey.
During an active survey, you'll first connect the Wi-Fi analyzer to the access points to measure the network strength and look for any weak spots. Then you'll walk through the entire location again to thoroughly test every inch. You can also generate heat maps again to visually represent problem areas.
You also need to check for channel overlapping. If neighboring access points are set to overlapping channels, you'll end up dealing with performance issues and network drops as devices move between the different WAPs.
The last type of site survey is a predictive survey. Predictive surveys use software programs to load the building blueprints and determines where to install the WAPs. This can be done remotely and doesn't require someone to walk through the location to test everything. A predictive survey isn't as precise as a passive or active survey, but it can cost less and take less time.

That's it for this lesson. In this lesson, we talked about the process of installing a wireless network. Before we can do anything, we need to know the usage needs of the network. Then you can perform passive, active, and predictive site surveys to determine how to install your WAPs.

8.4.3 Wireless Networking Facts

Wireless networking is commonplace, both in home and business environments.

This lesson covers the following topics:

- Wireless network hardware
- Wireless access point (WAP) placement
- Site surveys and heat maps
- Wi-Fi authentication methods
- Advanced authentication

Wireless Network Hardware

The following table describes hardware used in wireless networks:

Wireless Network Hardware	Description
Wireless access point (WAP)	A wireless access point broadcasts information and data over radio waves. A wireless access point functions as a wireless hub. The wireless access point may provide a connection to a physical wired network. The two classes of wireless access points are fat and thin. Fat access points have everything necessary to manage wireless clients and broadcast the network. Fat access points are standalone devices. Thin access points are basically a radio and antenna. Thin access points can broadcast the network, but require another system to manage clients and the network. Thin access points are referred to as controller-based devices. A WAP uses an service set identifier (SSID) that associates a name with a wireless network. This makes it easier for users to connect wirelessly.
Wireless interface	The wireless interface in a device, such as a laptop or smart phone, connects to the wireless access point.

Wireless bridge	A wireless bridge connects two wireless networks together. The bridge is typically created using a wired connection between the two access points. A bridge can be implemented wirelessly using a wireless distribution system (WDS).
Wireless LAN controller (WLC)	A Wireless LAN controller is used in a enterprise environment to manage multiple access points. The WLC is placed in the networking closet and connected to a switch. The controller is able to communicate with and manage the wireless access points. The WLC is also able to manage client connects and access point loads. This allows the WAPs to operate and work together as a single system instead of each device working in isolation.
Lightweight access point (LWAP)	Lightweight access points are used in conjunction with the wireless controller. LWAPs contain very little technology and rely on the WLC to handle everything including client connections, authentication, updating configurations, etc. LWAP forwards frames to the WLC to make the decision to either drop the packet or forward it. If the packet is to be forwarded, the WLC sends it to the applicable LWAP to which the destination device is connected and then that WLC sends the packet to the destination.

Wireless Access Point (WAP) Placement

Wireless network installation considerations refer to the factors that ensure good availability of authorized Wi-Fi access points. A network with patchy coverage is vulnerable to rogue and evil twin attacks.

The 5 GHz band has more space to configure non-overlapping channels. Also note that a WAP can use bonded channels to improve bandwidth, but this increases risks from interference.

An infrastructure-based wireless network comprises one or more wireless access points, each connected to a wired network. The access points forward traffic to and from the wired switched network. Each WAP is identified by its MAC address, also referred to as its basic service set identifier (BSSID). Each wireless network is identified by its name or service set identifier (SSID) .

Wireless networks can operate in either the 2.4 GHz or 5 GHz radio band. Each radio band is divided into a number of channels, and each WAP must be configured to use a specific channel. For performance reasons, the channels chosen should be as widely spaced as possible to reduce interference.

Site Surveys and Heat Maps

The coverage and interference factors mean that WAPs must be positioned and configured to cover the whole area with the least overlap as possible. A site survey is used to measure signal strength and channel usage throughout the area to cover. A site survey starts with an architectural map of the site, with features that can cause background interference marked. These features include solid walls, reflective surfaces, motors, microwave ovens, and so on. A Wi-Fi-enabled laptop or mobile device with Wi-Fi analyzer software installed performs the survey . The Wi-Fi analyzer records information about the signal obtained at regularly spaced points as the surveyor moves around the area.

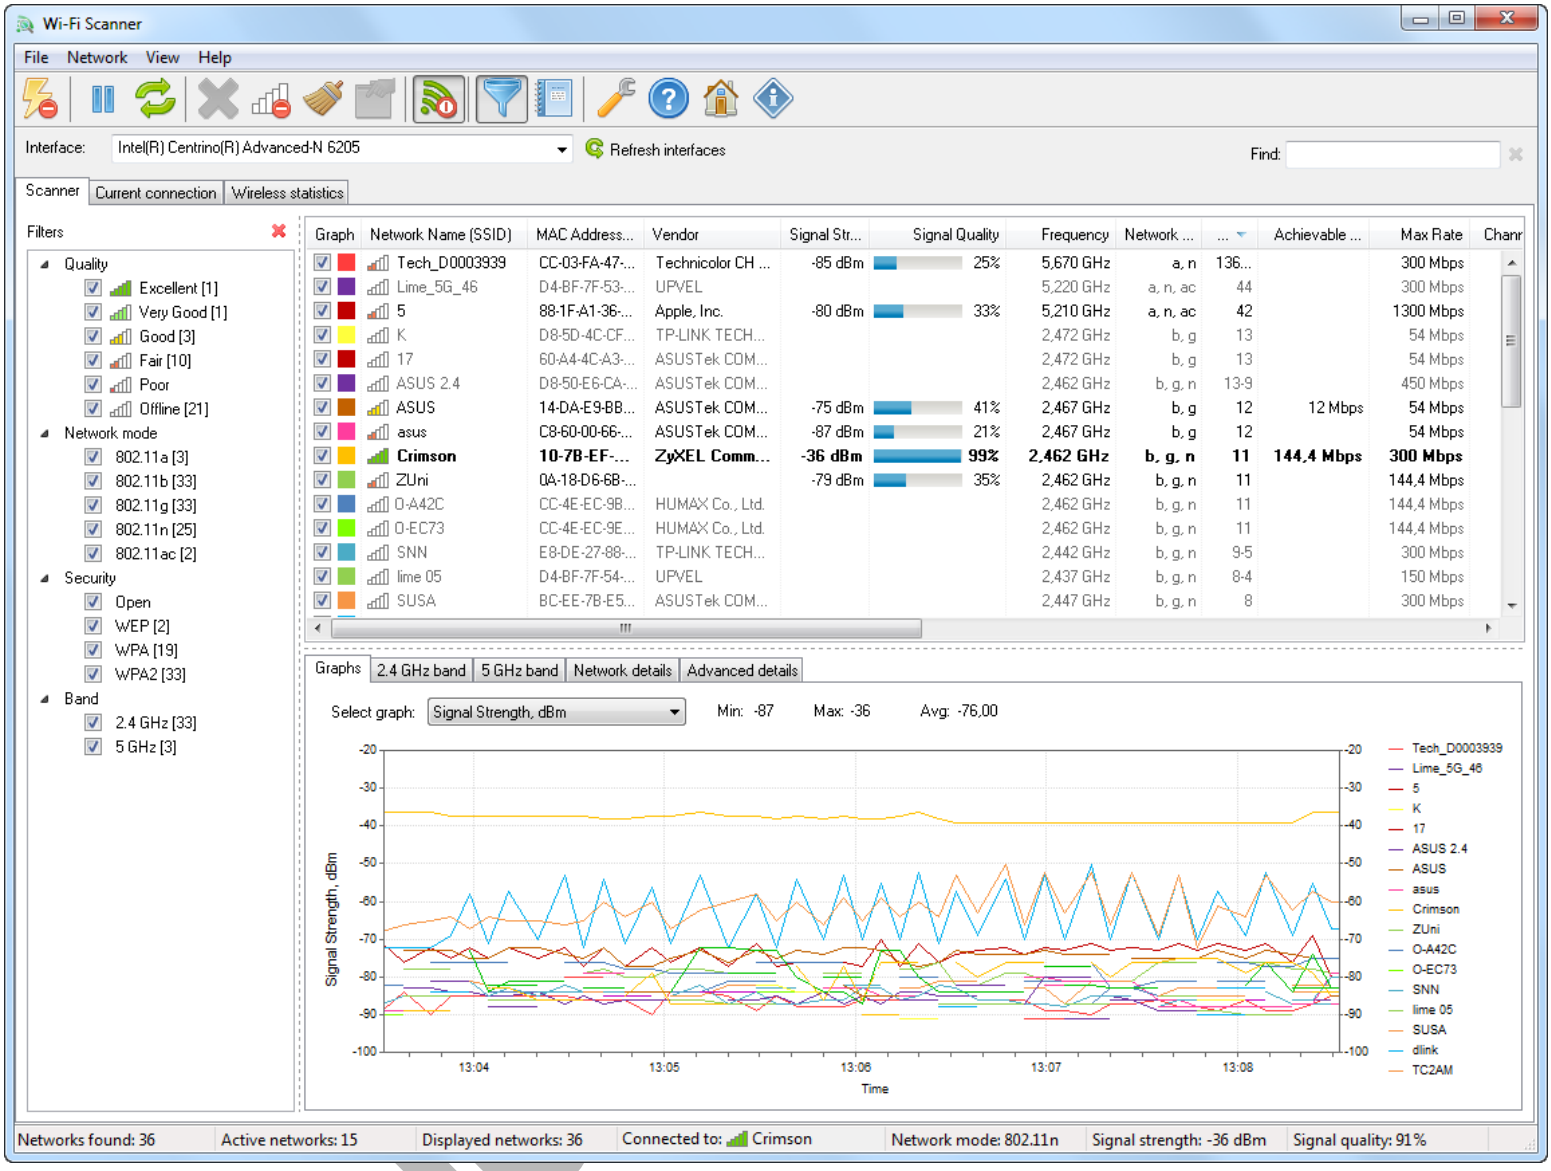

Example output from Lizard System's Wi-Fi Scanner tool. (Screenshot courtesy of Lizard Systems.)

These readings are combined and analyzed to produce a heat map , showing where a signal is strong (green/blue) or weak (red), and which channel is being used and how they overlap. This data is then used to optimize the design by adjusting transmit power to reduce a WAP's range, changing the channel on a WAP, adding a new WAP, or physically moving a WAP to a new location.

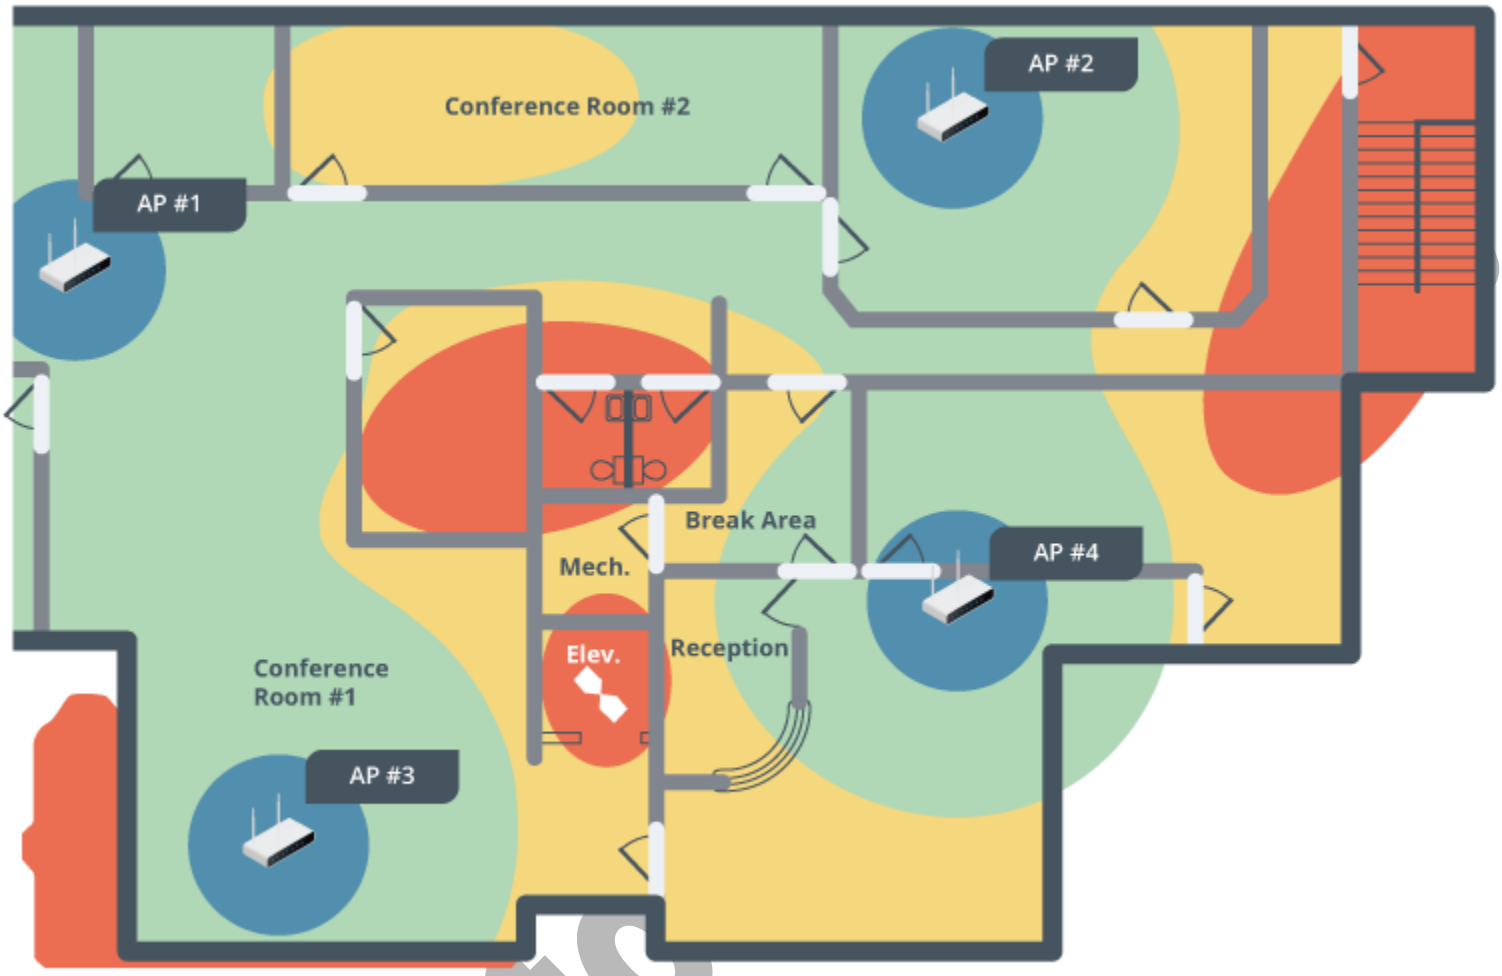

An illustration of a heat map.

As well as the site design, a wireless network must be configured with security settings. Without encryption, anyone within range can intercept and read packets passing over the wireless network. Security choices are determined by device support for the various Wi-Fi security standards, by the type of authentication infrastructure, and by the purpose of the WLAN. Security standard determine which cryptographic protocols are supported, the means of generating the encryption key, and the available methods for authenticating wireless stations when they try to join (or associate with) the network.

The first version of Wi-Fi Protected Access (WPA) was designed to fix critical vulnerabilities in the earlier wired equivalent privacy (WEP) standard. Like WEP, version 1 of WPA uses the RC4 stream cipher but adds a mechanism called the Temporal Key Integrity Protocol (TKIP) to make it stronger.

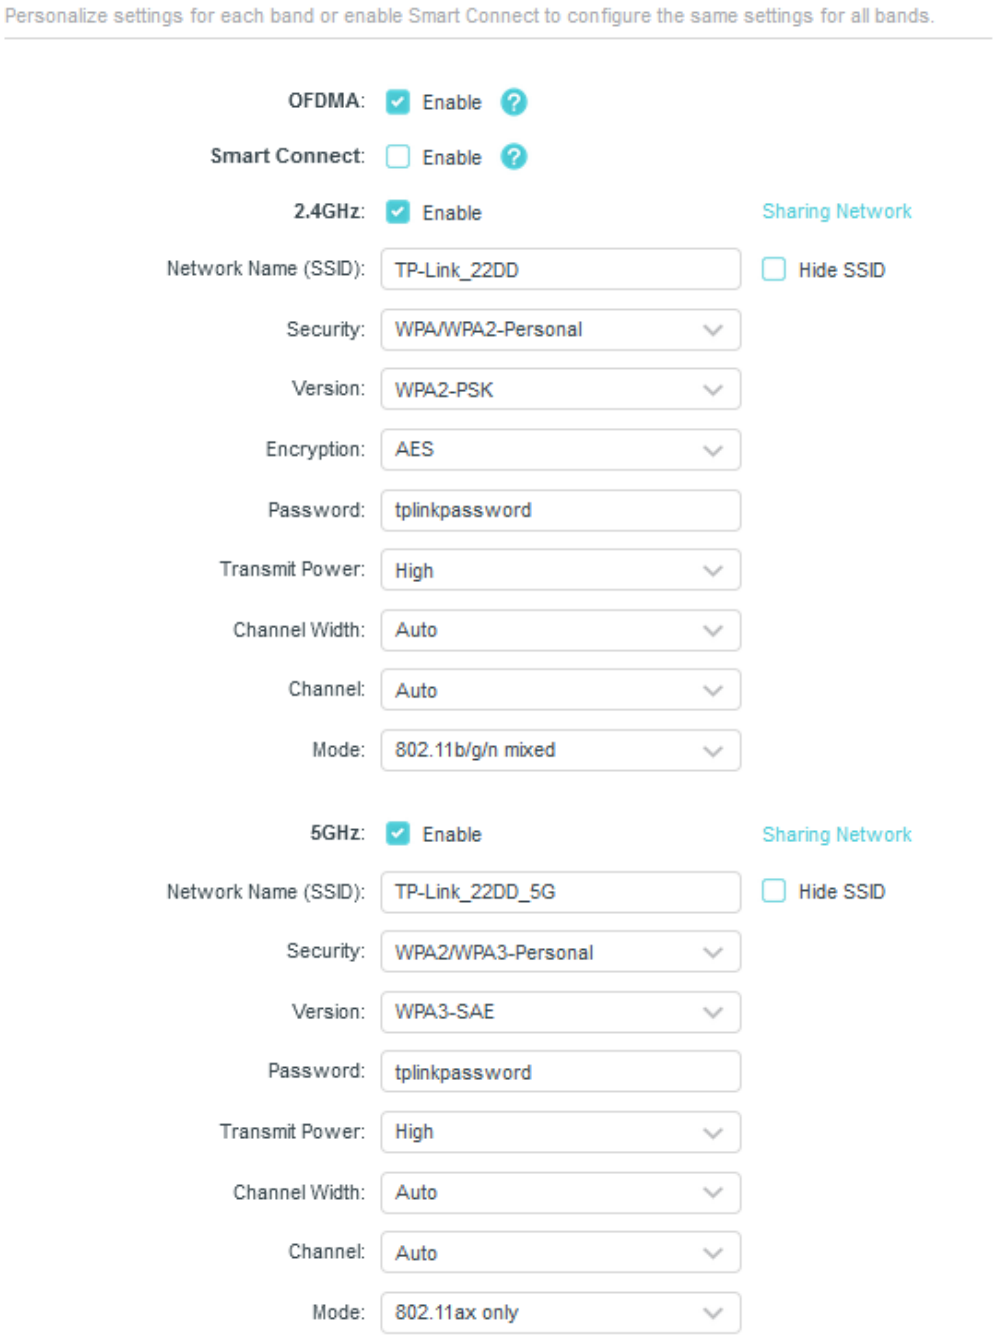

Configuring a TP-LINK SOHO access point with wireless encryption and authentication settings. In this example, the 2.4 GHz band allows legacy connections with WPA2-Personal security, while the 5 GHz network is for 802.11ax (Wi-Fi 6) capable devices using WPA3-SAE authentication. (Screenshot used with permission from TP-Link Technologies.)

As setting up an access point securely is relatively complex for residential consumers, vendors have developed a system to automate the process called Wi-Fi Protected Setup (WPS) . To use WPS, both the access point and wireless station (client device) must be WPS-capable. Typically, the devices will have a push button. Activating this on the access point and the adapter simultaneously will associate the devices using a PIN, then associate the adapter with the access point using WPA2. The system generates a random SSID and PSK. If the devices do not support the push button method, the PIN (printed on the WAP) can be entered manually.

Unfortunately, WPS is vulnerable to a brute force attack. While the PIN is eight characters, one digit is a checksum and the rest are verified as two separate PINs of four and three characters. These separate PINs are many orders of magnitude simpler to brute force, typically requiring just hours to crack. On some models, disabling WPS through the admin interface does not actually disable the protocol, or there is no option to disable it. Some APs can lock out an intruder if a brute force attack is detected, but in some cases, the attack can just be resumed when the lockout period expires.

To counter this, the lockout period can be increased. However, this can leave APs vulnerable to a denial of service (DoS) attack. When provisioning a WAP, it is essential to verify what steps the manufacturer has taken to make their WPS implementation secure and to use the required device firmware level identified as secure.

The Easy Connect method, announced alongside WPA3, is intended to replace WPS as a method of securely configuring client devices with the information required to access a Wi-Fi network. Easy Connect is a brand name for the Device Provisioning Protocol (DPP).

Each participating device must be configured with a public/private key pair. Easy Connect uses quick response (QR) codes or near-field communication (NFC) tags to communicate each device's public key. A smartphone is registered as an Easy Connect configurator app and associated with the WAP using its QR code. Each client device can then be associated by scanning its QR code or NFC tag in the configurator app. As well as fixing the security problems associated with WPS, this is a straightforward means of configuring headless Internet of Things (IoT) devices with Wi-Fi connectivity.

Wi-Fi Authentication Methods

In order to secure a network, you must confirm that only valid users are connecting to it. Wi-Fi authentication comes in three types: personal, open, and enterprise. Within the personal category, there are two methods: pre-shared key authentication (PSK) and simultaneous authentication of equals (SAE).

WPA2 Pre-Shared Key Authentication

In WPA2, pre-shared key (PSK) authentication uses a passphrase to generate the key used to encrypt communications. It is also referred to as group authentication because a group of users shares the same secret. When the access point is set to WPA2-PSK mode, the administrator configures a passphrase of between 8 and 63 ASCII characters. This is converted to a 256-bit HMAC (expressed as a 64-character hex value) using the PBKDF2 key stretching algorithm. This HMAC is referred to as the pairwise master key (PMK). The same secret must be configured on the access point and on each node that joins the network. The PMK is used as part of WPA2's 4-way handshake to derive various session keys.

All types of Wi-Fi personal authentication have been shown to be vulnerable to attacks that allow dictionary or brute force attacks against the passphrase. At a minimum, the passphrase must be at least 14 characters long to try to mitigate risks from cracking.

Wi-Fi Protected Access 3 (WPA3)

Neither WEP nor the original WPA version is considered secure enough for continued use. WPA2 uses the Advanced Encryption Standard (AES) cipher with 128-bit keys, deployed within the Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP). AES replaces RC4 and CCMP replaces TKIP. CCMP provides authenticated encryption, which is designed to make replay attacks harder.

Weaknesses found in WPA2 led to its intended replacement by WPA3. The main features of WPA3 are as follows:

- **Simultaneous Authentication of Equals (SAE)** - replaces the Pre-Shared Key (PSK) exchange protocol in WPA2, ensuring an attacker cannot intercept the Wi-Fi password even when capturing data from a successful login.

Enhanced Open - encrypts traffic between devices and the access point, even without a password, which increases privacy and security on open networks.

Updated Cryptographic Protocols - replaces AES CCMP with the AES Galois Counter Mode Protocol (GCMP) mode of operation. Enterprise authentication methods must use 192-bit AES, while personal authentication can use either 128-bit or 192-bit.

Wi-Fi Easy Connect - allows connecting devices by scanning a QR code, reducing the need for complicated configurations while maintaining secure connections.

Wi-Fi performance also depends on device support for the latest 802.11 standards. The most recent generation (802.11ax) is being marketed as Wi-Fi 6. The earlier standards are retroactively named Wi-Fi 5 (802.11ac) and Wi-Fi 4 (802.11n). The performance standards are developed in parallel with the WPA security specifications. Most Wi-Fi 6 devices and some Wi-Fi 5 and Wi-Fi 4 products should support WPA3 either natively or with a firmware/driver update.

While WPA3 still uses a passphrase to authenticate stations in personal mode, it changes the method this secret is uses to agree session keys. The scheme used is called a Password-Authenticated Key Exchange (PAKE). In WPA3, the Simultaneous Authentication of Equals (SAE) protocol replaces the 4-way handshake, which has been found vulnerable to various attacks. SAE uses the Dragonfly handshake, which is basically Diffie-Hellman over elliptic curves key agreement, combined with a hash value derived from the password and device MAC address to authenticate the nodes. With SAE, there should be no way for an attacker to sniff out the handshake to obtain the hash value and try to use an offline brute force or dictionary attack to recover the password. Dragonfly also implements ephemeral session keys providing forward secrecy.

The configuration interfaces for access points can use different labels for these methods. You might see WPA2-Personal and WPA3-SAE rather than WPA2-PSK and WPA3-Personal, for example. Additionally, an access point can be configured for WPA3 only or with support for legacy WPA2 (WPA3-Personal Transition mode). Researchers already found flaws in WPA3-Personal, one of which relies on a downgrade attack to use WPA2 (wi-fi.org/security-update-april-2019).

Advanced Authentication

Wireless enterprise authentication modes, such as WPA2/WPA3-Enterprise, include several essential components designed to improve security for corporate wireless networks. One important element is 802.1x authentication, which provides a port-based network access control framework, ensuring that only authenticated devices are granted network access. Typically, 802.1x requires an authentication server such as RADIUS (Remote Authentication Dial-In User Service), which verifies the credentials of users or devices trying to connect to the network.

In enterprise mode authentication schemes, users have a unique set of credentials rather than a shared passphrase as used in WPA2/WPA3 personal mode. Requiring each user or device to authenticate using unique credentials allows network administrators to track network usage at a granular level. The protocol also supports multiple Extensible Authentication Protocol (EAP) types, such as EAP-TLS, EAP-TTLS, or PEAP, which define specific authentication methods. EAP-TLS, for instance, uses client-server certificates for mutual authentication, while EAP-TTLS and PEAP utilize a server-side certificate. The server-side certificate is used to establish a secure tunnel for transmitting user credentials and helps devices validate the legitimacy of the access point. Enterprise mode authentication includes dynamic encryption key management, automatically changing the encryption keys used during a user's session.

Remote Authentication Dial-In User Service (RADIUS)

The Remote Authentication Dial-In User Service (RADIUS) standard is published as an Internet standard. There are several RADIUS server and client products.

The NAS device (RADIUS client) is configured with the IP address of the RADIUS server and with a shared secret. This allows the client to authenticate to the server. Remember that the client is the access device (switch, access point, or VPN gateway), not the user's PC or laptop. A generic RADIUS authentication workflow proceeds as follows:

The user's device (the supplicant) makes a connection to the NAS appliance, such as an access point, switch, or remote access server.

The NAS prompts the user for their authentication credentials. RADIUS supports PAP, CHAP, and EAP. Most implementations now use EAP, as PAP and CHAP are not secure. If EAP credentials are required, the NAS enables the supplicant to transmit EAP over LAN (EAPoL) data, but not any other type of network traffic.

The supplicant submits the credentials as EAPoL data. The RADIUS client uses this information to create an Access-Request RADIUS packet, encrypted using the shared secret. It sends the Access-Request to the AAA server using UDP on port 1812 (by default).

The AAA server decrypts the Access-Request using the shared secret. If the Access-Request cannot be decrypted (because the shared secret is not correctly configured, for instance), the server does not respond.

With EAP, there will be an exchange of Access-Challenge and Access-Request packets as the authentication method is set up and the credentials verified. The NAS acts as a pass-thru, taking RADIUS messages from the server, and encapsulating them as EAPoL to transmit to the supplicant.

At the end of this exchange, if the supplicant is authenticated, the AAA server responds with an Access-Accept packet; otherwise, an Access-Reject packet is returned.

Optionally, the NAS can use RADIUS for accounting (logging). Accounting uses port 1813. The accounting server can be different from the authentication server.

8.4.4 Configuring a Wireless Connection (Demo Video)

Transcript:

Connecting to Wi-Fi is something pretty much everyone does these days, but there are some additional settings that many users are not aware of. In this demonstration, we are going to go and see where you can configure these additional settings.

I have a PC that has a USB Wi-Fi adapter. When click on the taskbar you can see we are not connected to any networks. Our Wi-Fi adapter is on however its just not connected to anything. When we click this arrow, we are shown the available wireless networks around this area.
To connect to a wireless network, I find the SSID, which is right here. I can click it once and I'm able to click Connect. From here, I'm prompted for the passphrase. We have that set up to be LearningIsCool. I can then click here to verify that I typed it in correctly and click Next. Great were connected and ready to go.
Let's look at our wireless profiles. To see these, we need to go to **Start** >**Settings**>Once you successfully connect to a wireless, the computer creates what's called a profile.
The profile identifies wireless networks that have been connected to it in the past. These profiles are saved so you can connect to the same network again later.
Click on the profile to change its properties about the connection.
The first thing I can do is tell this to connect automatically to this wireless. If it's one that you want to use as your default, you might want to have this turned to On.
If you pay for bandwidth, you might want to tell this connection that you're on a metered connection. This will help prevent things such as Windows Update from downloading large files while connected. This is nice because you can then avoid charges for data you didn't plan on using. Let's click up here and go back.
You might want to add a wireless network profile manually. This might be if you're connecting to a wireless network and there is no SSID broadcast.

To do this, we select Add a new network. We can then enter our network name here. I'll put in Testout_Hidden and for Security type, I'll select WPA2-Personal from the list. I need to enter a security key which will be the same LearningIsCool. I then have the option to connect automatically and also to connect even if this network isn't broadcasting. I'll click Save here at the bottom, and we've added a network manually. Since were already connected to an existing network we will have to go back down to your Wi-Fi connected and click connect on the Testout_Hidden. Once doing so it will connect to this network now. Keep note you may only want to update what Wi-Fi networks you want to connect to automatically if there are many around the same location.
To remove a network you no longer want, you select the Wi-Fi connection that you want to remove and Right Click to select the Forget option. This doesn't show you a confirmation dialog, so pay attention and make sure you're removing the right one. This can also be done in the Manage known networks.
Mobile Hotspot is a feature that allows you to create a hotspot on your computer. Instead of you connecting to a hotspot you are making this computer the hotspot. An example could be having one computer plugged into a wired internet connection and then the same computer provided wireless to share the internet connection. Not only can it make a hotspot off a wired connection it can reshare a wireless connection too. In our task bar there is a quick option to turn on mobile hotspot. After clicking on it our mobile hotspot will be active. In order to see more settings for this we must go to or Network& internet settings. Were already in our settings so we will take a look at it from here. The setting were looking for is Mobile hotspot. The settings here provide key information such as the password for this Mobile hotspot. If we select edit were able to change the Network name and password if need be. If you're on your battery this setting here for Power saving will be helpful because it will automatically shut Mobile hotspot off once it notices no devices are connecting to it. The last thing I would like to point out is this hotspot is only limited to 8 devices so if you have more devices then that then this setup will not work for you.
A lot of the available features for a wireless adapter are still available however it may be hidden for the public eye. A regular user may not be concerned about them, but we are. Open the search and type cmd. This will bring up the option to open the command prompt. In our command prompt let's type in "netsh wlan show wirelesscapabilities". Press enter. This list will show us what is supported by the wireless adapter and what is not. For example, ANQP Service Information Discovery was the previous option listed in windows 10 for hotspot 2.0. its not listed in our settings on this Windows install however its still supported by the wireless adapter.
That's it for this demo. In this demo, we looked at how to connect to a wireless network in Windows. We then looked at Mobile hotspot and supported capabilities of a wireless adapter.

8.4.5 Configure a Wireless Network (Simulation)

Scenario

You are a network technician for a small corporate network. You just installed a Ruckus zone controller and wireless access points throughout your office buildings using wired connections. You now need to configure basic wireless network settings.

Access the Wireless Controller console through Chrome on **http://192.168.0.6** with the username **admin** and the password **password** . The username and password are case-sensitive.

In this lab, your task is to:

Create a WLAN using the following settings:

Name: **CorpNet Wireless**

ESSID: **CorpNet**

Type: **Standard Usage**

Authentication: **Open**

Encryption: **WPA2**

Encryption algorithm: **AES**

Passphrase: **@CorpNetWeRSecure!**

Connect the Exec-Laptop in the Executive office to the new wireless network.

Explanation

Complete this lab as follows:

Access the Ruckus zone controller.

From the taskbar, select **Google Chrome** .

In the URL field, enter **192.168.0.6** and press **Enter** .

Maximize the window for better viewing.

Log into the Wireless Controller console.

In the Admin field, enter **admin** (case-sensitive).

In the Password field, enter **password** as the password.

Select **Login** .

Create a new WLAN.

Select the **Configure** tab.

From the left menu, select **WLANs** .

From the right, under WLANs, select **Create New** .

In the New Name field, enter **CorpNet Wireless** .

In the ESSID field, enter **CorpNet** .

Under Type, make sure **Standard Usage** is selected.

Under Authentication Options, make sure **Open** is selected.

Under Encryption Options, select **WPA2** .

For Algorithm, make sure **AES** is selected.

In the Passphrase field, enter **@CorpNetWeRSecure!** .

Select **OK** .

Switch to the Exec-Laptop.

From the top left, select **Floor 1** .

Under Executive Office, select **Exec-Laptop** .

Connect to the new CorpNet wireless network.

In the notification area, select the ***wireless network*** icon to view the available networks.

Select **CorpNet** .

Select **Connect** .

Enter **@CorpNetWeRSecure!** for the security key.

Select **Next** .

Select **Yes** to make the computer discoverable on the network. The CorpNet network now shows as being connected and secured.

8.4.6 Practice Questions (Section Quiz)

q_wireless_access_point_secp8

Which of the following is used on a wireless network to identify the network name?

Answers:

***SSID**

MAC address

IP address

Subnet mask

Explanation:

Wireless devices use the service set identifier (SSID) to identify a network name. All devices on a wireless network use the same SSID.

The MAC address is a unique physical device address.

The IP address is a logical address that includes both the logical network and the logical device address.

The subnet mask is used with the IP address to identify the network portion of the IP address.

q_wireless_heat_map_secp8

Which of the following is generated after a site survey and shows the Wi-Fi signal strength throughout the building?

Answers:

***Heat map**

Diagram

Analyzer

Ad hoc

Explanation:

A heat map is generated following a site survey. A heat map shows the Wi-Fi signal strength in different locations.

A diagram of the location is needed so survey results can be overlaid.

A Wi-Fi analyzer is used to perform a site survey.

Ad hoc wireless configuration mode provides wireless communication without a wireless access point. This is not a type of site survey.

q_wireless_pake_secp8

You are a network security consultant for a small business that is setting up a new wireless network. The business owner is concerned about the security of the network, especially in terms of protecting against unauthorized access.

The owner is not technically savvy and wants a solution that is secure but also user-friendly.

Which of the following authentication methods would you recommend?

Answers:

Pre-shared key (PSK)

Extensible Authentication Protocol (EAP)

***Password-Authenticated Key Exchange (PAKE)**

Open System Authentication

Explanation:

Password-Authenticated Key Exchange (PAKE) is the correct answer. It is a method that allows users to establish a secure communication channel by using a password. It provides a good balance between security and user-friendliness. Even if an attacker intercepts the password, they cannot use it to derive the encryption key. This makes PAKE a good choice for a small business owner who wants a secure but user-friendly authentication method.

While PSK is relatively user-friendly, it is not the most secure option. If the pre-shared key is compromised, the entire network is at risk. Furthermore, managing PSKs can become complex as the number of users increases.

Extensible Authentication Protocol (EAP) is a secure authentication framework that supports multiple authentication methods. However, it can be complex to set up and manage, especially for a small business owner who is not technically savvy.

Open System Authentication offers no real authentication or security. It simply allows any device to connect to the network as long as it knows the network's SSID. This would not be a suitable choice for a business owner who is concerned about network security.

q_wireless_placement_secp8

The IT manager has tasked you with installing the new Wireless LAN controller (WLC).

Where should you install the controller?

Answers:

Lobby

***Network closet**

Manager's office

Roof

Explanation:

A WLC should be placed in the networking closet and connected to a switch so it can communicate with and manage the wireless access points.

None of the other locations are valid locations to install the WLC.

q_wireless_radius_auth_secp8

You are a network administrator for a large organization that uses RADIUS for network authentication. One day, you receive a report that a user is unable to access the network.

You check the RADIUS server logs and see that the user's authentication request was received, but no Access-Accept or Access-Reject message was sent back to the user.

Which of the following could be the MOST likely cause of this issue?

Answers:

The user entered the wrong password.

***The RADIUS server is down.**

The user is not in the RADIUS server's database.

The network connection between the user and the RADIUS server is unstable.

Explanation:

The RADIUS server is down is the correct answer. If the RADIUS server is down, it will not be able to send any messages back to the user. This would explain why no Access-Accept or Access-Reject message was sent back to the user.

If the user entered the wrong password, the RADIUS server would still send an Access-Reject message back to the user. Therefore, this is not the most likely cause of the issue.

If the user is not in the RADIUS server's database, the server will still send an Access-Reject message back to the user. Therefore, this is not the most likely cause of the issue.

If the network connection was unstable, the RADIUS server might not have received the user's authentication request in the first place. However, the logs show that the server did receive the request. Therefore, this is not the most likely cause of the issue.

q_wireless_security_01_secp8

Which type of wireless access point is generally used in a residential setting?

Answers:

LWAP

***SOHO**

Bridge

WLC

Explanation:

In a small office or residential location, a Small Office Home Office (SOHO) wireless router is often used. These devices are three different devices in one:

A router function connects the internal LAN to the internet.

A switch portion connects the internal wired LAN devices.

An access point portion allows the internal wireless devices to connect to the network.

Lightweight access points (LWAPs) are used in conjunction with a wireless controller.

A wireless bridge connects two wireless networks.

A Wireless LAN controller (WLC) is used in an enterprise environment to manage multiple access points.

q_wireless_security_02_secp8

You need to implement a solution to manage multiple access points in your organization.

Which of the following would you MOST likely use?

Answers:

LWAP

SOHO

Bridge

***WLC**

Explanation:

A Wireless LAN controller (WLC) is used in an enterprise environment to manage multiple access points. A WLC is placed in the networking closet and connected to a switch. The controller is able to communicate with and manage the wireless access points.

In a small office or residential location, a Small Office Home Office (SOHO) wireless router is often used.

Lightweight access points (LWAPs) are used in conjunction with a wireless controller.

A wireless bridge connects two wireless networks.

q_wireless_site_survey_01_secp8

Which of the following devices would you use to perform a site survey?

Answers:

Heat map

Wireless access point

***Wi-Fi analyzer**

Wireless interface

Explanation:

A Wi-Fi analyzer is used to perform a site survey. A Wi-Fi analyzer can be a specialized tool or a software program running on a laptop, smartphone, or tablet.

A heat map is generated following a site survey. A heat map shows the Wi-Fi signal strength in different locations.

A wireless access point (WAP) broadcasts information and data over radio waves. WAPs function as wireless hubs.

A wireless interface in a device, such as a laptop or smartphone, connects to a wireless access point.

q_wireless_site_survey_02_secp8

Which of the following types of site surveys should be performed first?

Answers:

Active

***Passive**

Predictive

Ad hoc

Explanation:

An initial site survey performed should be a passive survey. This survey is performed without the analyzer connecting to any specific WAP and is in a listen-only mode.

An active survey is performed after multiple passive surveys have been completed and the wireless access points have been placed. An active survey verifies proper coverage has been achieved.

A predictive survey uses software programs to load the building blueprints and determine where to install the WAPs.

An ad hoc wireless configuration mode provides wireless communication without a wireless access point. Ad hoc mode is not a type of site survey.

q_wireless_wap_01_secp8

Which of the following is responsible for broadcasting information and data over radio waves?

Answers:

***Wireless access point**

Wireless bridge

Wireless interface

Wireless LAN controller

Explanation:

A wireless access point (WAP) broadcasts information and data over radio waves. WAPs function as wireless hubs.

A wireless bridge connects two wireless networks.

A wireless interface in a device, such as a laptop or smartphone, connects to a wireless access point.

A Wireless LAN controller is used in an enterprise environment to manage multiple access points.

q_wireless_wap_02_secp8

Which class of wireless access point (WAP) has everything necessary to manage clients and broadcast a network already built into its functionality?

Answers:

Thin

Bridge

Ad hoc

***Fat**

Explanation:

Fat access points have everything necessary to manage wireless clients and broadcast a network. Fat access points are standalone devices.

Thin access points are basically a radio and antenna. Thin access points can broadcast a network but require another system to manage clients and the network.

A wireless bridge connects two wireless networks.

Ad hoc wireless configuration mode provides wireless communication without a wireless access point.

q_wireless_wap_03_secp8

The network administrator of an educational institution is upgrading an existing wireless network. The campus has various buildings, each having multiple floors, and the aim is to ensure consistent Wi-Fi coverage across the entire campus.

To achieve this, a site survey and heat map creation will guide the placement and configuration of wireless access points (WAPs).

Which of the following would MOST accurately represent the correct actions based on the survey results?

Answers:

Place WAPs in areas indicated as red in the heat map and increase transmit power to the maximum in all devices.

***Place WAPs in areas with weak signals (green/blue in the heat map) and carefully manage transmit power to avoid unnecessary overlap.**

Ignore the heat map and place WAPs in every room regardless of signal strength to ensure maximum coverage.

Set all WAPs to use the same channel to ensure a seamless transition for users moving between different WAP coverage areas.

Explanation:

WAPs belong in areas where the signal is weak to ensure consistent coverage. Managing the transmit power helps to fine-tune the coverage area, reducing unnecessary overlap and interference.

Placing WAPs in areas already marked red (strong signal) on the heat map would lead to unnecessary overlap, causing interference.

This approach could lead to too much overlap and interference between WAPs, resulting in poor performance. It also ignores the valuable information the heat map provides.

Setting all WAPs to the same channel can cause co-channel interference. Instead, WAPs should use different, non-overlapping channels to avoid interference. The channel selection should follow the site survey results to optimize the network performance.

q_wireless_wap_04_secp8

The IT department of a medium-sized company is preparing to implement a new wireless network to enhance mobility and productivity within the workplace.

The IT team has conducted an in-depth analysis of the office layout and user requirements and now faces the task of strategically positioning the wireless access points (WAPs) throughout the office. The team understands the potential risks and challenges involved in wireless deployments.

Which considerations should the IT team take into account when strategically positioning the wireless access points (WAPs) in the company's new wireless network? (Select two.)

Answers:

***Optimal coverage of the office area to ensure reliable connectivity for all users.**

***Minimizing interference from nearby electronic devices to enhance network performance.**

Positioning WAPs based on the design preferences of the company's management.

Ensuring the WAPs are accessible only to employees with specific job roles and responsibilities.

Making sure the service set identifier (SSID) that associates a name with a wireless network is disabled.

Explanation:

The following are considerations that the IT team should take into account:

Strategically positioning the wireless access points (WAPs) to achieve optimal coverage is essential to ensure reliable connectivity for all users throughout the office and eliminate potential dead zones in every corner of the workplace.

Minimizing interference from nearby electronic devices is crucial for enhancing network performance and reducing connectivity issues. To mitigate this, the IT team should identify potential sources of interference and strategically place the WAPs away from these devices.

Network performance and coverage should take precedence over design preferences to ensure optimal connectivity and productivity.

Access control measures determine who can connect to the wireless network and what level of privileges they have. While important for network security, this consideration is distinct from the physical placement of WAPs.

A WAP uses a service set identifier (SSID) that associates a name with a wireless network. This makes it easier for users to connect wirelessly. Disabling the SSID would disable the WAP service.

q_wireless_wpa3_upgrade_secp8

You are a network administrator for a large corporation that handles sensitive data. The corporation is currently using WPA2 for Wi-Fi authentication.

However, there have been recent concerns about the security of the network. You have been tasked with deciding whether to upgrade the Wi-Fi authentication method.

Which of the following would be the BEST choice?

Answers:

Continue using WPA2.

Switch to WEP.

Switch to WPA.

***Upgrade to WPA3.**

Explanation:

Upgrade to WPA3 is the correct answer. WPA3 is the latest and most secure Wi-Fi authentication method currently available. It includes several improvements over WPA2, such as stronger encryption and more secure password-based authentication. Therefore, upgrading to WPA3 would be the best choice for a corporation handling sensitive data.

While WPA2 is more secure than its predecessor, WPA, it has known vulnerabilities that can be exploited by determined attackers. Therefore, continuing to use WPA2 would not be the best choice for a corporation handling sensitive data.

WEP is an older encryption standard and is considered to be the least secure. It has numerous well-known vulnerabilities and would not be a suitable choice for a corporation handling sensitive data.

WPA is more secure than WEP but less secure than WPA2 and WPA3. It was only intended as an interim solution during the development of WPA2. Therefore, switching to WPA would not be the best choice.

8.5 Wireless Attacks

As you study this section, answer the following questions:

What is the difference between bluejacking and bluesnarfing?

What is an initialization vector used for?

How can you discover rogue access points?

What is the difference between passive and active radio frequency identification (RFID) tags?

In this section, you will learn to:

Detect rogue hosts.

Configure rogue host protection.

The key terms for this section include:

Term	Definition
Rogue access points	Any unauthorized access point added to a network.

Initialization vector (IV)	A seed value used in encryption. The seed value and the key are used in an encryption algorithm to generate additional keys or encrypt data.
Radio frequency identification (RFID)	RFID uses radio waves to transmit data from small circuit boards called RFID tags to special scanners.
Near Field Communication (NFC)	NFC allows two-way communication between two devices. The devices must be within a few centimeters of each other. NFC is a newer technology built on RFID.
Interference	A signal that corrupts or destroys a wireless signal. Interference can affect communication of access points and other wireless devices.
Evil twin	A wireless access point that deceives users into believing that it is a legitimate network access point.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.2 Explain common threat vectors and attack surfaces. Unsecure networks Wireless Bluetooth 2.4 Given a scenario, analyze indicators of malicious activity. Physical attacks Radio frequency identification (RFID) cloning Network attacks Wireless Credential replay 4.1 Given a scenario, apply common security techniques to computing resources Wireless devices

	Installation considerations Wireless security settings Wi-Fi Protected Access 3 (WPA3) Authentication protocols
TestOut Security Pro	2.2 Harden network devices 2.2.2 Configure and access a wireless network 2.2.4 Harden a wireless network

8.5.1 Wireless Attacks (Lesson Video)

Transcript:

Because wireless networks use an unbounded communication method, they're especially vulnerable to attacks. This includes all types of wireless networks, as in Wi-Fi, Bluetooth, RFID, and even NFC. In this lesson, I'm going to go over the various attacks that threaten wireless networks.

Because network security is a top priority, it's common for network administrators to implement stringent rules and limit website access. Some users find this very frustrating.
To try to evade these rules, users sometimes install an access point on their own computer. This is known as a soft access point. A network hacker who has gained access to a network, perhaps using MAC spoofing, might do the same. This is called a rogue access point or an unauthorized association.
Here's the big problemâ€”most employees don't understand the importance of securing their soft access point or even how to do it. So they just leave it open to hackers. When an attacker discovers an access point, they're able to run various types of vulnerability scanners from outside the company, perhaps from a car or adjacent building or even from several miles away. Keep in mind that if a hacker can gain some type of physical access to your company, they can also hide a physical rogue access point as well. This is often done by configuring an extremely compact and powerful hardware device called a Raspberry Pi as their access point.
A rogue access point that's placed on the network by a malicious attacker can be used to run what's called an evil twin attack. Here's how this worksâ€”say your organization currently has a wireless network with a SSID of MyCompanyWiFi. The attacker configures their own access point with the same SSID and places it near the organization's building. Then the attacker uses a jamming or disassociation attack to knock users off the legitimate network. When users reconnect to the network, they're now connecting to the attacker's access point instead.
Once a victim is connected to a rogue access point, the attacker can capture all data flowing through it. The user shouldn't notice anything different since their internet is still running like normal. These attacks are extremely dangerous as the attacker gains immediate access to all sorts of sensitive information.
To defend against these attacks, you should perform regular site surveys to locate any rogue access points. Your wireless networks encrypt data, which helps protect against sniffing. But some older standards, like WEP, are vulnerable due to improperly configured initialization vectors, or IVs.
When we encrypt data, we use a secret key that only we know. If we use the same key for everything we encrypt, it won't take long for the key to be compromised. This is where initialization vectors can help out. The IV is a random string that's added to the key to create a different key for each encryption process. Unfortunately, if we don't configure the IV properly, it can open us up to attacks.
Wired Equivalent Privacy, or WEP, was one of the first wireless security protocols that used IVs to help with encryption. The problem, though, was the IV was only 24 bits long and the key was 40 bits. This meant there was only approximately 16 million different IVs that could be generated. This may sound like a lot, but it actually meant that an IV would be repeated at least once every 4096 packets. It doesn't take much time to get 4096 packets, and attackers were able to develop programs that would flood the network with packets to find matching IVs. Once an attacker gets multiple IVs, they can begin decrypting the encryption key. After that, all communication is open and no longer secure.

No networks should use WEP anymore. Newer wireless standards such as WPA-2 don't even use IVs for encryption. For many of these wireless attacks to work, the attacker needs to kick everyone off the legitimate network and prevent them from reconnecting. They attempt this with jamming attacks or disassociation attacks.

A Wi-Fi jamming attack requires a special device that works by analyzing the spectrum used by wireless networks and then transmitting a powerful signal on the same frequencies. Basically, the jammer tries to be the loudest voice in the room so nearby devices can't see the legitimate wireless network. The hope is that they disconnect and are unable to reconnect again. Jamming devices are illegal and difficult to come by, though.

A disassociation attack, on the other hand, can be performed with a laptop. This is also know as a deauthentication attack. When a device connects to a wireless network, special unencrypted management packets are sent back and forth. A deauthentication attack takes advantage of this unencrypted process by sending fake, malicious deauthentication packets to kick people off the network. The attacker can select individual users or kick everyone off at once. Jamming and deauthentication attacks have the same end result, but they use very different methods.

Wi-Fi is used everywhere today and attackers are always working on methods to exploit weaknesses. Yet another wireless technology that can be susceptible to attacks is Bluetooth.

From headphones and speakers to watches, fitness devices, automobiles, and even refrigerators, Bluetooth devices are heavily integrated into our daily lives. This makes these devices attractive to malicious attackers. An attacker can gather all sorts of personal data and even take over our devices' embedded cameras and microphones.

One common Bluetooth attack is Bluejacking. When performing a Bluejacking attack, he or she looks for another Bluetooth device that's discoverable and sends unwanted messages to it. Attackers can't gain control of a device or steal any data so Bluejacking is more annoying than harmful.

Bluesnarfing, on the other hand, can be harmful. Bluesnarfing is when an attacker exploits a vulnerability in the OBject EXchange Protocol, or OBEX Protocol, to pair to a target device. Once paired, he or she is able to access all the victim's device data and then disconnect without leaving a trace.

The best defense against both of these attacks is to only enable Bluetooth when you need to and to not leave devices in Discovery Mode.

Radio frequency identification, or RFID, uses radio waves to transmit data from small circuit boards called RFID tags to special scanners. Active RFID tags, like those seen in a car's toll pass, have on-board batteries and can send signals over a long distance. Passive RFID tags, which are used in ID badges and credit cards, aren't powered. These tags instead rely on the RFID scanner's transferred energy.

A malicious attacker can build their own scanner to target these RFID tags. If an attacker can get close enough, they can scan the RFID chip and clone it. An attacker might also eavesdrop on the communication between the chip and reader and steal the data that way.

To protect against these attacks, RFID chips often operate at different frequencies, making it more difficult for an attacker to quickly complete a scan. Also, because most of these chips are passive, an attacker has to get extremely close for it to work. Near field communication, or NFC, is a newer technology that's built on RFID. NFC allows two-way communication between two devices. But they must be within a few centimeters of each other.

NFC is what allows us to use our smartphones to pay for items at the store. We can also use NFC to transfer files between two phones if they're right next to each other. NFC is vulnerable to the same types of attacks as RFID, but it requires the attacker to be incredibly close.

That'll wrap up this lesson. In this lesson, we looked at some different wireless communications attacks. We first looked at Wi-Fi and rogue access points, evil twin attacks, how an IV can be exploited, and also how jamming and disassociation attacks work. Then we looked at Bluetooth attacks, including Bluejacking and Bluesnarfing. Finally, we covered how RFID and NFC work and the security concerns associated with both of these devices.

8.5.2 Wireless Attack Facts

This lesson covers the following topics:

- Wireless attacks
- Rogue access points
- RFID/NFC attacks

Wireless Attacks

Wireless networks present particular security challenges and are frequently the vector for various types of attacks. Many organizations have set up and configured wireless networks, and if not configured properly, these networks can be susceptible to attack.

The following table explains common wireless attacks:

Wi-Fi Attack	Description
Wireless denial-of-service	A wireless denial-of-service (DoS) attack is usually designed to prevent clients from connecting to the legitimate access point. A wireless network can be disrupted by interference from other radio sources. These are often unintentional, but it is also possible for an attacker to purposefully jam the legitimate network by setting up a rogue access point with a stronger signal. Wireless DoS can also target clients. In the normal course of operations, an access point and client use management frames to control connections. A disassociation attack exploits the lack of encryption in management frame traffic to send spoofed frames. One type of disassociation attack injects management frames that spoof the MAC address of a single victim station in a disassociation notification, causing it to be disconnected from the network. Another variant of the attack broadcasts spoofed frames to disconnect all stations. As well as trying to redirect connections to an evil twin, a disassociation attack might also be used in conjunction with a replay attack aimed at recovering the network key.
Wireless replay and key recovery	Wireless authentication is vulnerable to various types of replay attacks that aim to capture the hashes used when a wireless station associates with an access point. Once the hash is captured, it can be subjected to offline brute force and dictionary cracking. A KRACK attack uses a replay mechanism that targets the WPA and WPA2 4-way handshake. KRACK is effective regardless of whether the authentication mechanism is personal or enterprise. It is important to ensure both clients and access points are fully patched against such attacks.
Evil twin attack	Rogue APs placed by an attacker can be used to run an evil twin attack. In this attack: The rogue AP is configured to mimic the legitimate network. The attacker uses a jamming or disassociation attack to knock users off the legitimate network. When users reconnect to the network, they connect to the attacker's AP. The attacker can monitor and capture all data that moves through the rogue AP. To protect against this attack, conduct a radio frequency (RF) noise analysis to detect a malicious rogue AP that uses jamming to force wireless clients to connect to it instead of legitimate APs.

Initialization vector (IV) attack	An initialization vector is a seed value used in encryption. The seed value and the key are used in an encryption algorithm to generate additional keys or to encrypt data. Wired Equivalent Privacy (WEP) encryption reuses initialization vectors. The reuse of IVs makes it easy for attackers to crack them. This is known as an IV attack. Be aware that: The WEP IV is 24-bits, and the key is 40-bits. This allows for approximately 16 million IVs. An IV is repeated at least once every 4096 packets. Hackers developed programs that flood the network with packets, allowing them to find matching IVs quickly. Once enough IVs are obtained, the attacker can decrypt the encryption key. WEP encryption can be cracked in as little as 1-2 minutes. Due to the vulnerabilities of WEP, you should no longer use it. Newer standards, such as WPA2 and WPA3, do not use IVs in the encryption process.
Jamming attack	With wireless networks, interference is a signal that corrupts or destroys the wireless signal sent by APs and other wireless devices. Non-malicious interference includes the following: Electromagnetic interference (EMI) is interference caused by motors, heavy machinery, and fluorescent lights. Radiofrequency interference (RFI) is interference on the radio channel. It is caused by nearby wireless devices using the same channel, cordless phones, or microwave ovens. Adjacent channels on wireless APs have a small degree of overlap. To avoid interference with other wireless APs within the same vicinity, use channels that don't overlap neighboring wireless APs. Some interference is malicious in nature and designed to disrupt wireless network communications. Malicious interference is sometimes referred to as jamming. In a jamming attack, a transmitter is tuned to the same frequency and the same type of modulation as the wireless network. The jamming signal overrides the legitimate wireless network radio signals at the receiving devices. The following list describes different types of jamming signals that can be used to disrupt a Wi-Fi network. *Spark jamming* is the most effective type of Wi-Fi interference attack. It repeatedly blasts receiving equipment with high-intensity, short-duration RF (radio frequency) bursts at a rapid pace. Experienced

	RF signal technicians can usually identify this type of attack quickly because of the regular nature of the signal. *Random noise jamming* produces radio signals using random amplitudes and frequencies. While not as effective as a spark attack, the random noise attack is harder to identify due to the intermittent and random nature of the interference. In fact, this type of signal is frequently mistaken for background radio noise that occurs naturally. *Random pulse jamming* uses radio signal pulses of random amplitude and frequency to interfere with a Wi-Fi network.
Disassociation/deauthentication attack	Wireless devices are vulnerable to deauthentication (deauth) and disassociation attacks because the 802.11 standard allows devices to be authenticated with multiple APs at once. When a device connects to a wireless network, special unencrypted management packets are sent back and forth. Deauthentication and disassociation attacks take advantage of these packets to disconnect devices from a network. Be aware that: To execute a deauth attack, the attacker pretends to be the wireless router the device is connected to. The attacker disconnects the device from the network. When the user tries to reconnect, the attacker can intercept the user's information. Disassociation attacks are similar. Instead of disconnecting a user, disassociation tricks the user into giving the fake router responsibility for forwarding packets.

Rogue Access Points

A rogue access point is one that has been installed on the network without authorization, whether with malicious intent or not. A malicious user can set up such an access point with something as basic as a smartphone with tethering capabilities, and a non-malicious user could enable such an access point by accident. If connected to a local segment, an unauthorized access point creates a backdoor through which to attack the network.

A rogue access point masquerading as a legitimate one is called an evil twin. Each network is identified to users by a service set identifier (SSID) name. An evil twin might use typosquatting or SSID stripping to make the rogue network name appear similar to the legitimate one. Alternatively, the attacker might use some DoS technique to overcome the legitimate access point. In the latter case, they could spoof both the SSID and the basic SSID (BSSID). The BSSID is the MAC address of the access point's radio. The evil twin might be able to harvest authentication information from users entering their credentials by mistake and implement a variety of other on-path attacks, including DNS redirection.

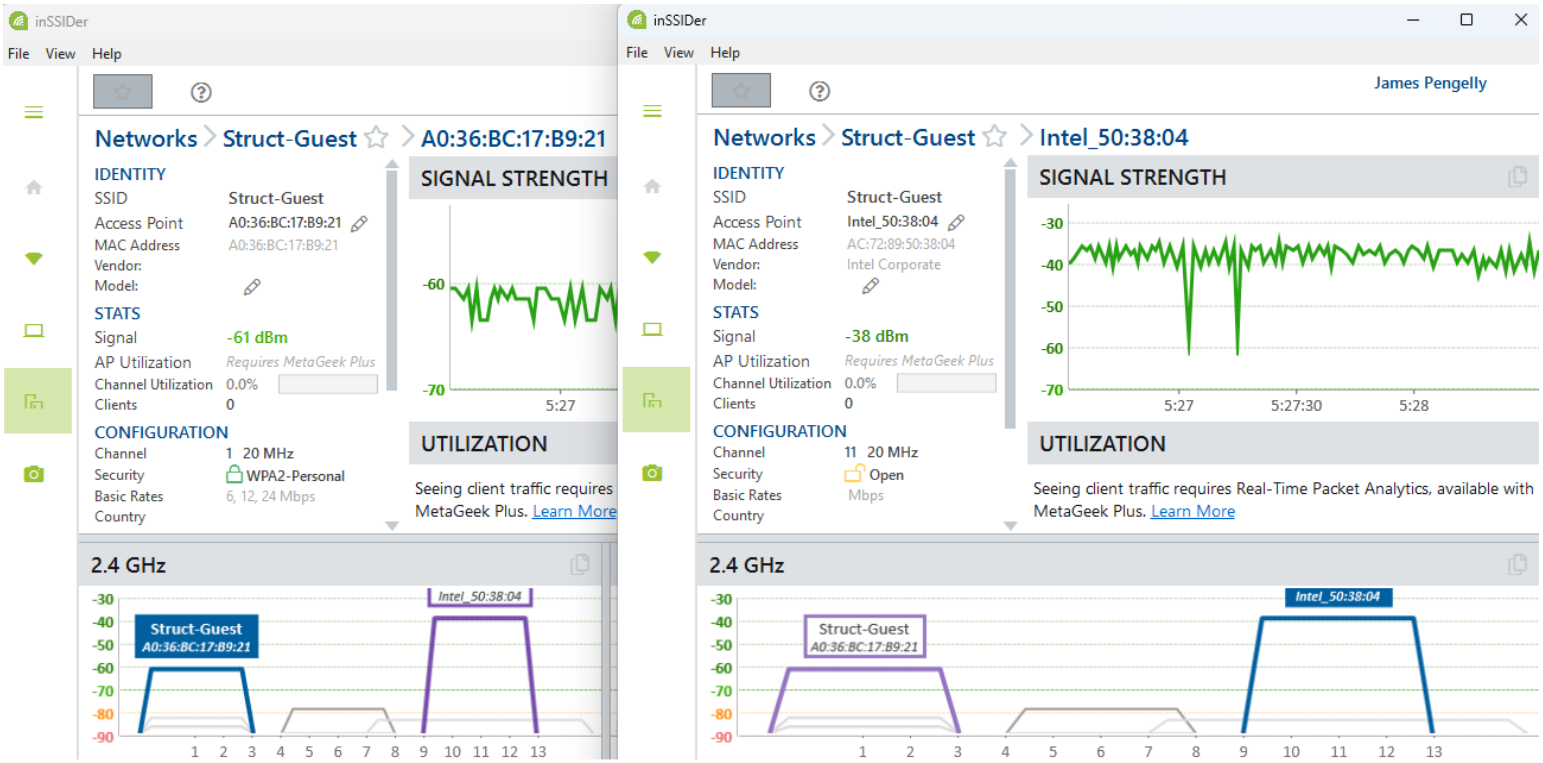

Surveying Wi-Fi networks using MetaGeek inSSIDer. The Struct-Guest network shown in the first window is the legitimate one and has WPA2 security configured. The evil twin has the same SSID but a different BSSID (MAC address) and open authentication. (MetaGeek, Inc. © Copyright 2005-2023.)

A rogue hardware access point can be identified through physical inspections. There are also various Wi-Fi analyzers and wireless intrusion protection systems that can detect rogue access points. These can log the use of typosquatting SSIDs and unknown and duplicate (spoofed) MAC addresses. In an enterprise network, access points are usually connected to switches. Monitoring can detect any that are not and flag them as potential rogues. They may also be able to identify radio hardware and alert if an unauthorized access point brand is detected.

Bluetooth Attacks

Bluetooth is designed to allow devices to communicate within a personal area network (PAN) of close proximity. PAN devices include cell phones, personal digital assistants (PDAs), printers, mice, and keyboards.

Bluetooth:

- Is designed for distances longer than infrared (IR) communication and has lower power consumption.
- Requires that devices be in discovery mode to find each other and synchronize.
- Operates in the 2.4 GHz range and uses adaptive frequency hopping (AFH).

Eavesdropping is difficult because Bluetooth implements authentication and key derivation with custom algorithms based on the SAFER+ block cipher. It also uses the E0 stream cipher for encrypting packets. Bluetooth is one of the most secure protocols for mobile device communication, but it is susceptible to the following attacks:

Bluejacking looks for nearby devices that are in discovery mode and sends unwanted messages. The attacker is unable to steal any data. This attack is more annoying than harmful.

Bluesnarfing exploits a vulnerability in the object exchange (OBEX) protocol that allows an attacker to pair to the target device. Once paired, the attacker can view the calendar, emails, text messages, contact lists, and other data on the device. Many Bluetooth devices have built-in features to prevent bluesnarfing, but it is still a known vulnerability.

To mitigate the risks of Bluetooth attacks, enable Bluetooth only when needed and make sure discovery mode is turned off except when pairing devices.

RFID/NFC Attacks

RFID uses radio waves to transmit data from small circuit boards, called RFID tags, to special scanners. There are two types of RFID tags:

Active RFID tags have onboard batteries and can send signals over a long distance. Road toll passes and other types of passes use active RFID.

Passive RFID is not powered and relies on the energy of the scanner to transmit data. These tags are seen in ID badges, credit cards, and similar devices.

RFID systems are vulnerable to various kinds of attacks, including:

RFID Attack	Description
Eavesdropping	An attacker uses an RFID reader to listen to conversations between a tag and the intended reader.
Man-in-the-middle (MTM)	An attacker intercepts a signal from an RFID tag and then manipulates it before sending it to the intended recipient. This kind of attack is frequently used to take down a system.
Denial-of-service (DOS)	An attacker blocks radio signals or jams the system with interfering noise.
Cloning and spoofing	An attacker creates a copy of an existing tag and uses the fake tag to gain access to a secure system.

To protect against these attacks, RFID chips often operate at different frequencies. This makes it more difficult for an attacker to find and scan them.

Near Field Communication (NFC) is a newer technology that is built on RFID. NFC allows two-way communication between two devices. The devices must be within a few centimeters of each other. Although NFC transmission distances are very short, transmissions are susceptible to several malicious attacks, including:

A lost NFC device allows anyone who finds it to access NFC resources.

NFC signals can be jammed by malicious interference.

NFC devices and readers are susceptible to man-in-the-middle exploits, where an attacker captures transmissions from the reader and forwards them onto the device, potentially reading and modifying data in transit.

NFC devices and readers are susceptible to relay attacks. An attacker can capture NFC data in transit and use the information to masquerade as the original device.

8.5.3 Detecting Rogue Hosts (Demo Video)

Transcript:

In this demo, we're going to configure our wireless controller to detect rogue hosts and look at other features wireless controllers have. Some of these features help with rogue host detection. If you remember, a rogue host is an unauthorized access point within your network or a different network in the surrounding area.

Let's start by logging in to our controller. On the main page, we have some statistics that can be useful for seeing what's going on with your wireless networks. We're going to check to see if this controller is set up for detecting rogue hosts. Let's go to Insight > Untrusted Rogue APs. Notice there are no APs listed in this section. To populate this list, we're going to have to trigger a scan to look for these them.
Under Access Points, we're going to click the one access point we have. This first window shows details related to the access point. Now we're going to go to Configuration and see what kind of settings we have available for this AP.
IP Setting allows you to set DHCP or static addresses. If settings are set to DHCP, you can specify a failback IP address, MASK, and gateway. Radio gives you several options to tweak, such as mode, channel width, channel, and TX power. Some controllers have an auto mode for Tx power, but this one doesn't.
Load balance provides options for balancing connected clients. Max associated clients can be set so once that threshold is hit, load balancing is triggered. RSSI Threshold is a signal strength setting. If the client's signal is weaker than the set RSSI threshold, it'll disconnect the client from the AP. WLANs show which wireless networks are set up with this access point. You can override each access point if you don't want the same SSIDs on each access point. LED is just the indicator light on the access point.
The last setting we're going to look at is Rogue AP Detection. This is the scan button we use to scan for Rogue APs. Let's start a scan. Depending on how many surrounding networks there are, it could take several minutes. Great, the scan has completed. Let's go take a look at the list.
We're going to go back to insight and then to the Untrusted Rouge APs. This will show us some details related to information gathered during the scan. Let filter them so we know which APs are the closest by signal strength. You can see we have one AP that's close with -46 signal strength.
Sometimes, it's even possible to see the device's location on a map with other solutions. We have two options for rogue APs: we can trust them, or we can delete them from the list. The Tp-link Omada controller already knows that it's an untrusted AP, and it'll remain that way unless we trust it.
Let's say we know about this device, so we'll mark it as trusted. The device was removed from the untrusted list, and now it's on the Trusted APs list.
There could be a time where you may want to remove a trusted rouge AP and put it back on the untrusted list. To do that, you click the Untrust button. Now the SSID network NotYourWifi is back on the untrusted list.
It's important to know that this rogue device detection feature is just one step in protecting your network. If you were to identify a very suspect rogue device here, you would probably want to take additional steps, such as using an RF scanner to pinpoint the device and physically remove it from your organization or find the building where the network is coming from. Not all rogue APs are expected to be on your network, as other users may set up APs for their own use.
That's it for this demo. In this demo, we discussed the configuration of various features on a wireless controller. We looked at various settings we can control with an AP. We also looked at how you can manage rogue devices.

8.5.4 Configure Rogue Host Protection (Simulation)

Scenario

You are a network technician for a small corporate network. You want to take advantage of the self-healing features provided by the small enterprise wireless solution you've implemented. You're already logged in as WxAdmin on the Wireless Controller console from ITAdmin.

In this lab, your task is to:

Configure self-healing on the wireless network.

Automatically adjust AP radio power to optimize coverage when interference is present.

Set 2.4 GHz and 5 GHz radio channels to use the **Background Scanning** method to adjust for interference.

Configure the background scanning needed for rogue device detection, AP locationing, and self-healing. Background scans should be performed on all radios every **30 seconds** .

Configure load balancing for all radios by adjusting the threshold to **40 dB** .

Configure band balancing to allow no more than **30%** of clients to use the 2.4 GHz radios.

Reduce the power levels to **-3 dB** for three access points in Building A to reduce RF emanations. Use the wireless survey results in the exhibit to identify the access points.

The amount you reduce TX Power by requires a judgment call based on the wireless survey results. In practice, you would repeat the wireless survey to verify the proper TX Power settings.

Explanation

Complete this lab as follows:

Configure self-healing.

From the Ruckus ZoneDirector, select the **Configure** tab.

From the left menu, select **Services** .

From the right, under *Self-Healing,* select **Automatically adjust AP radio power to optimize coverage when interference is present** .

Using the *Automatically adjust 2.4GHz channels using* drop-down menu, select **Background Scanning** from the drop-down menu.

Using the *Automatically adjust 5GHz channels using* drop-down menu, select **Background Scanning** from the drop-down menu.

From the right of this section, select **Apply** .

Configure background scanning.

Under Background Scanning, select **Run a background scan on 2.4GHz radio** .

Enter **30** seconds.

Select **Run a background scan on 5GHz radio** .

Enter **30** seconds.

From the right of this section, select **Apply** .

Configure load balancing.

Under Client Load Balancing, select **Run load balancing on 2.4GHz radio** .

In the *Adjacent radio threshold(dB)* field, enter **40** .

Select **Run load balancing on 5GHz radio** .

In the *Adjacent radio threshold(dB)* field, enter **40** .

From the right of this section, select **Apply** .

Configure band balancing.

Under Band Balancing, select **Percent of clients on 2.4GHz radio** .

Enter **30** .

From the right of this section, select **Apply** .

Adjust the AP power level.

From the left menu, select **Access Points** .

From the top right, select **Exhibit** to determine which access points to adjust.

Select **Edit** next to the access point to be modified.

Under *Radio B/G/N(2.4G)* next to TX Power, make sure **Override Group Config** is selected.

From the *TX Power* drop-down list, select **-3dB (1/2)** .

Under *Radio A/N/AC(5G)* next to **TX Power** , make sure **Override Group Config** is selected.

From the *TX Power* drop-down list, select **-3dB (1/2)** .

Select **OK** .

Repeat steps 5b - 5h for additional access points.

8.5.5 Practice Questions (Section Quiz)

q_wl_attacks_bluejacking_secp8

Which of the following sends unsolicited business cards and messages to a Bluetooth device?

Answers:

***Bluejacking**

Bluesnarfing

Bluebugging

Slamming

Explanation:

Bluejacking is a relatively harmless practice that entails an unknown sender sending business cards anonymously to a Bluetooth recipient within a distance of 10-100 meters, depending on the class of the Bluetooth device. The business cards usually include a flirtatious message so the attacker can see a visual reaction from the recipient. Multiple messages are sent to the device if the attacker thinks there is a chance they will be added as a contact. Bluetooth devices are not susceptible to bluejacking if they are set to non-discoverable mode.

Bluesnarfing is the use of a Bluetooth connection to gain unauthorized access to an existing Bluetooth connection between phones, desktops, laptops, or PDAs. Bluesnarfing allows the attacker to view calendars, emails, text messages, and contact lists.

Bluebugging gives an attacker access to all mobile phone commands that use Bluetooth technology, such as initiating phone calls, sending and receiving messages, eavesdropping, and reading and writing phone book contacts.

Slamming entails unauthorized or fraudulent changes made to a subscriber's telephone service or DSL internet service.

q_wl_attacks_bluesnarfing_secp8

Which of the following BEST describes bluesnarfing?

Answers:

Sending anonymous electronic business cards

***Viewing calendar, emails, and messages on a mobile device without authorization**

Executing commands on a mobile device

Cloning a mobile device

Explanation:

Bluesnarfing is the use of a Bluetooth connection to gain unauthorized access to an existing Bluetooth connection between phones, desktops, laptops, or PDAs. Bluesnarfing allows access to view the calendar, emails, text messages, and contact lists. Many Bluetooth devices have built-in features to prevent bluesnarfing, but it is still a known vulnerability.

Bluejacking is a relatively harmless practice that entails an unknown sender sending business cards anonymously to a Bluetooth recipient within a distance of 10-100 meters, depending on the class of the Bluetooth device. The business cards usually include a flirtatious message so the attacker can see a visual reaction from the recipient. Multiple messages are sent to the device if the attacker thinks there is a chance they will be added as a contact. Bluetooth devices are not susceptible to bluejacking if they are set to non-discoverable mode.

Bluebugging gives an attacker access to all mobile phone commands that use Bluetooth technology, such as initiating phone calls, sending and receiving messages, eavesdropping, and reading and writing phone book contacts. Only highly skilled individuals can perform bluebugging.

q_wl_attacks_bluetooth_secp8

A large hospital uses Bluetooth technology for short-range personal area networking. The organization has a security concern with bluesnarfing.

What is the attacker doing to the organization?

Answers:

***Stealing information from someone else's phone by using an exploit in Bluetooth.**

Sending unsolicited text messages or vCards to a discoverable device.

Launching highly effective attacks using a peripheral device with malicious firmware.

Compromising any active and unpatched system, regardless of whether discovery is enabled.

Explanation:

Using an exploit in Bluetooth, bluesnarfing allows attackers to steal information from someone else's phone, circumventing the authentication mechanism and accessing sensitive information without the user's knowledge or consent.

Sending unsolicited text messages or vCards to a discoverable device is bluejacking, where an attacker sends an unsolicited message or contact details.

Connecting to malicious peripherals presents a risk of launching highly effective attacks using a peripheral device with malicious firmware. However, it is not specifically related to bluesnarfing.

The BlueBorne exploit compromises any active and unpatched system, regardless of whether the system enables discovery, compromising Bluetooth-enabled systems.

q_wl_attacks_emi_secp8

Which type of interference is caused by motors, heavy machinery, and fluorescent lights?

Answers:

RFI

NFC

***EMI**

RFID

Explanation:

Electromagnetic interference (EMI) is interference caused by motors, heavy machinery, and fluorescent lights.

Radio frequency interference (RFI) is interference on the radio channel. It is caused by nearby wireless devices using the same channel, cordless phones, or microwave ovens.

Near frequency communication (NFC) allows two-way communication between two devices. The devices must be within a few centimeters of each other.

Radio frequency identification (RFID) uses radio waves to transmit data from small circuit boards, called RFID tags, to special scanners.

q_wl_attacks_evil_twin_secp8

Which of the following BEST describes an evil twin?

Answers:

***An access point configured to mimic a valid access point to obtain logon credentials and other sensitive information.**

A threat agent that marks the outside of buildings to indicate the presence of a wireless network.

An access point that is added to a network by an internal employee to provide unauthorized network access.

A Bluetooth device that receives mobile phone commands via bluebugging.

Explanation:

An evil twin is a rogue access point configured to mimic a valid access point. In contrast, a rogue access point is any unauthorized access point added to a network. The evil twin may be configured to prompt for credentials, allowing the attacker to steal those credentials or use them in a man-in-the-middle attack to connect to the valid wireless access point.

Warchalking is marking the outside of buildings to indicate the presence of a wireless network. Attackers might use these marks to alert others of open or secured wireless networks. Businesses might even use these marks to advertise free wireless networks.

Bluebugging gives an attacker access to all mobile phone commands that use Bluetooth technology, such as initiating phone calls, sending and receiving messages, eavesdropping, and reading and writing phone book contacts. Only highly skilled individuals can perform bluebugging.

q_wl_attacks_iv_secp8

Which type of attack is WEP extremely vulnerable to?

Answers:

***IV attack**

Evil twin

Bluesnarfing

Cloning

Explanation:

Wired Equivalent Privacy (WEP) is extremely vulnerable to initialization vector (IV) attacks because WEP reuses the IVs. This makes it easy for attackers to crack them and compromise the encryption.

An evil twin attack is a type of rogue access point attack.

Bluesnarfing is a Bluetooth attack.

Cloning is an RFID attack.

q_wl_attacks_jamming_secp8

You are the security analyst for your organization. Clients are complaining about being unable to connect to the wireless network. After looking into the issue, you have noticed short bursts of high-intensity RF signals are interfering with your wireless network's signal.

Which type of attack are you MOST likely experiencing?

Answers:

Disassociation

***Jamming**

Bluesnarfing

Cloning

Explanation:

In a jamming attack, a transmitter is tuned to the same frequency and type of modulation as the wireless network. The jamming signal overrides the legitimate wireless network radio signals. This scenario is a spark jamming attack.

A disassociation attack occurs when a user is tricked into giving a fake router responsibility for forwarding packets.

Bluesnarfing is a Bluetooth attack.

Cloning is an RFID attack.

q_wl_attacks_nfc_01_secp8

An attacker has intercepted near-field communication (NFC) data and is using that information to masquerade as the original device.

Which type of attack is being executed?

Answers:

Disassociation

Bluesnarfing

Cloning

***Relay**

Explanation:

This scenario describes a relay attack. A relay attack occurs when an attacker can capture NFC data in transit and use the information to masquerade as the original device.

A disassociation attack occurs when a user is tricked into giving a fake router responsibility for forwarding packets. This is not performed on NFC devices.

Bluesnarfing is a Bluetooth attack.

Cloning occurs when an attacker creates a copy of an existing RFID tag and uses the fake tag to gain access to a secure system.

q_wl_attacks_nfc_02_secp8

A cybersecurity analyst is implementing security measures for Near Field Communication (NFC) usage in the organization's mobile devices.

Which technique should the analyst consider applying to mitigate potential risks associated with NFC technology?

Answers:

Enable NFC chip reading for all devices to enhance connectivity options.

Use NFC for direct payment transactions without the need for mobile wallet apps.

***Apply encryption to NFC data to prevent eavesdropping and on-path attacks.**

Increase the NFC signal range to improve communication.

Explanation:

Applying encryption to NFC data is crucial to prevent eavesdropping and on-path attacks, ensuring that sensitive information remains secure during communication.

Enabling NFC chip reading for all devices may increase connectivity options, but it can also expose them to potential risks and unauthorized access.

q_wl_attacks_rfid_secp8

Which type of RFID tag can send a signal over a long distance?

Answers:

Passive

***Active**

NFC

Bluetooth

Explanation:

Active RFID tags have onboard batteries and can send signals long distances. Road toll passes and other types of passes use active RFID.

Passive RFID is not powered and relies on the energy of the scanner to transmit data. These tags are seen in ID badges, credit cards, and similar devices.

NFC allows two-way communication between two devices. The devices must be within a few centimeters of each other.

Bluetooth is designed to allow devices to communicate within a personal area network of close proximity.

q_wl_attacks_rogue_01_secp8

Your company security policy states that wireless networks are not to be used because of the potential security risk they present to your network.

One day, you find that an employee has connected a wireless access point to the network in his office.

Which type of security risk is this?

Answers:

***Rogue access point**

Man-in-the-middle attack

Phishing

Physical security

Social engineering

Explanation:

A rogue access point is an unauthorized access point added to a network, or it is an access point that is configured to mimic a valid access point. Examples include:

- An attacker or an employee with access to the wired network installs a wireless access point on a free port. The access port then provides a way to access the network remotely.
- An attacker near a valid wireless access point installs an access point with the same (or similar) SSID. The access point is configured to prompt for credentials, allowing the attacker to steal those credentials or use them in a man-in-the-middle attack to connect to the valid wireless access point.
- An attacker configures a wireless access point in a public location and then monitors traffic to see who connects to the access point.

A man-in-the-middle attack is used to intercept information passing between two communication partners. A rogue access point might be used to initiate a man-in-the-middle attack. But in this case, the rogue access point was connected without malicious intent. Social engineering exploits human nature by convincing someone to reveal information or perform an activity. Phishing uses an email and a spoofed website to gain sensitive information.

q_wl_attacks_rogue_02_secp8

A user contacts a company help desk complaining about intermittent connection problems to needed network files and shares. The user also noticed connection problems occur when the network signal strength is at its highest.

What could this be a sign of? (Select two.)

Answers:

***Rogue access point**

***Wireless denial-of-service**

Wireless replay

Downgrade attack

RFID attack

Explanation:

A rogue access point (AP) allows a person with malicious intent to place a rogue AP with a higher power to capture usernames and passwords without getting caught immediately.

A wireless denial-of-service attack causes network problems by not allowing users access to legitimate APs due to the rogue AP's higher signal strength.

A wireless replay attack aims to capture the hashes used when a wireless station associates with an access point. Then, the hash gets subjected to attacks to crack the encryption. It does not impact wireless access.

A downgrade attack attempts to make a server or client use a lower specification protocol with weaker ciphers and key lengths. It does not impact wireless access.

RFID uses radio waves to transmit data from small circuit boards, called RFID tags, to special scanners. It does not impact wireless access to network files and shares.

q_wl_attacks_site_survey_secp8

You are concerned that wireless access points may have been deployed within your organization without authorization.

What should you do? (Select two. Each response is a complete solution.)

Answers:

***Conduct a site survey.**

***Check the MAC addresses of devices connected to your wired switch.**

Implement an intrusion detection system (IDS).

Implement an intrusion prevention system (IPS).

Implement a network access control (NAC) solution.

Explanation:

A rogue host is an unauthorized system connected to a wireless network. It could be an unauthorized wireless device, or it could even be an unauthorized wireless access point that someone connected without permission to a wired network jack. Rogue hosts could be benign in nature, or they could be malicious. Either way, rogue hosts on your wireless network could represent a security risk and should be detected and removed if necessary. Four commonly used techniques for detecting rogue hosts include:

Using site survey tools to identify hosts and APs on the wireless network

Checking connected MAC addresses to identify unauthorized hosts

Conducting an RF noise analysis to detect a malicious rogue AP that is using jamming to force wireless clients to connect to it instead of legitimate APs

Analyzing wireless traffic to identify rogue hosts

Using an IDS or an IPS would not be effective, as these devices are designed to protect networks from perimeter attacks. Rogue APs are internal threats. A NAC solution can be used to remediate clients that connect to a network, but a NAC solution can't be used to detect a rogue AP.

q_wl_attacks_wpa3_secp8

A company's IT security specialist decides to upgrade the wireless network infrastructure to enhance data protection during transmissions. Recognizing the importance of strong encryption for wireless data, the specialist evaluates the various encryption standards available.

Which wireless encryption standard offers the MOST robust security for protecting wireless data transmissions and has become the preferred choice for many organizations?

Answers:

***WPA3**

WEP

WPA

TKIP

Explanation:

Wi-Fi Protected Access 3 (WPA3) enhances security beyond its predecessors and is the most robust wireless data protection encryption standard. It rectifies vulnerabilities found in WPA2 and WPA and represents the latest standard.

People view wired equivalent privacy (WEP) as a dated wireless encryption standard and deem it weak due to its identifiable vulnerabilities. It does not measure up to WPA3 in security.

TKIP is a WLAN encryption protocol designed to provide more secure encryption than WEP. However, it does not provide the same level of security as WPA3.

8.6 Wireless Defenses

As you study this section, answer the following questions:

- Which settings in a wireless access point can you configure to improve security?
- Which cryptographic protocol uses a Remote Authentication Dial-In User Service (RADIUS) server?
- Which access method forces a user to view and interact with it before accessing a network?
- What are the three components in an 802.1x setup?
- Which EAP standard is considered to be one of the most secure?

In this section, you will learn to:

- Harden a wireless network.
- Configure a wireless intrusion prevention system.

The key terms for this section include:

Term	Definition
Wi-Fi Protected Access (WPA)	The most commonly used cryptographic protocol in use for wireless networks. WPA2 and WPA3 are the two versions in use.
Pre-shared key (PSK)	Wireless access method that utilizes a passphrase for users to connect.
Wi-Fi Protected Setup (WPS)	Wireless access method that allows a device to securely connect to a wireless network without typing the PSK.
Open network	Wireless access method that has no authentication.
Captive portal	Wireless access method that forces a user to view and interact with it before accessing a network.
802.1x	Standard for local area networks that is used to authenticate users to a wireless network. It was created by The Institute of Electrical and Electronics Engineers Standards Association (IEEE-SA).
Remote Authentication Dial-In User Service (RADIUS)	A protocol used to authenticate users in a enterprise environment to a wireless network.

Extensible Authentication Protocol (EAP)	An authentication framework that uses a set of interface standards. EAP allows various authentication methods to be used.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.2 Explain common threat vectors and attack surfaces. Unsecure networks Wireless Default credentials 2.3 Explain various types of vulnerabilities. Hardware Firmware 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Hardening techniques Default password changes 3.2 Given a scenario, apply security principles to secure enterprise infrastructure. Infrastructure considerations Network appliances Intrusion prevention system (IPS)/intrusion detection system (IDS) Port security 802.1X Extensible Authentication Protocol (EAP) Secure communication/access Tunneling

	Transport Layer Security (TLS) 4.1 Given a scenario, apply common security techniques to computing resources. Wireless devices Installation considerations Wireless security settings Wi-Fi Protected Access 3 (WPA3) Cryptographic protocols AAA/Remote Authentication Dial-In User Service (RADIUS) Authentication protocols
TestOut Security Pro	2.2 Harden Network Devices 2.2.2 Configure and Access a Wireless Network 2.2.4 Harden a Wireless Network

8.6.1 Wireless Security (Lesson Video)

Transcript:

Wireless networking uses radio frequencies to transmit data. This means that anyone with a wireless receiver can capture data from a network if it's not secured properly. In this lesson, we'll go over some common configuration problems and different cryptographic protocols we can implement to further secure our data.

When we set up a new wireless access point, the first thing we need to do is properly configure it. Many users just plug it in and use the default settings, which is extremely dangerous. These settings are readily available for anyone to look up on the internet.
The first setting to change is the default login information. All access points are shipped with a default login, such as admin. No other setting matters more than this one because if an attacker has our login information, they basically have complete control of the whole device.
The next setting we should change is the network name, or SSID, which stands for service set identifier. The SSID can be a maximum of 32 bytes in length. We want to make sure that our SSID is unique to ourselves, but it shouldn't include identifiable information such as last name or address.
By default, the SSID is broadcast, and anyone within range can see the network name. We can, though, disable this right from the start if we choose. This might seem like an obvious setting to enable, but it doesn't actually provide as much security as you'd think. A determined hacker can easily discover hidden SSIDs, and beyond that, disabling the SSID can cause connection issues for your device. You should experiment to see if disabling your SSID is the best option for your network.
You can also configure your wireless access point to use MAC address filtering. This makes it so only hosts that have certain MAC addresses can connect, which helps you control which systems are allowed access in the first place. MAC address filtering provides a minimal level of security at best, though. It may discourage a casual attacker, but it won't stop a determined one.
If an attacker is really set on gaining access to your wireless network, they can sniff packets and see which MAC addresses are being used to connect to the wireless access point. Then they can use MAC address spoofing to configure their system with one of the allowed MAC addresses.
You should also monitor your wireless access point antenna placement and wireless access point antenna power levels.

Data emanation is a significant security problem. By default, the radio signals used by a wireless network broadcast omni-directionally. That radio signal travels farther than you realize, which makes it possible for an attacker outside of your building to access your wireless network. You can limit this exposure by manipulating your wireless access point antenna placement and your wireless access point power levels to reduce the amount of data emanation outside of the building.

Proper access point configuration is important to keeping a wireless network safe. But the most important security step is choosing the correct encryption method. Using a weak encryption method like WEP would make a network vulnerable to any hacker.

When we configure our wireless network, it's important that we select the best cryptographic protocols available. For most users, the Wi-Fi Protected Access, or WPA, versions 2 or 3 are the best options.

WPA2 was first introduced in 2004 and is still widely used today. There are two versions availableâ€”WPA2-Personal and WPA2-Enterprise. WPA2-Personal, also known as WPA2-PSK, or WPA2 pre-shared key, protects our network by using a pre-shared key referred to as the passphrase. WPA2-Enterprise verifies users through a RADIUS server.

WPA2-Personal uses AES with CCMP to encrypt all data. This stands for Advanced Encryption Standard with Counter Mode Cypher Block Chaining Message Authentication Code Protocol. AES with CCMP uses a 128-bit key and encrypts data in 128-bit blocks.

When a device connects to the access point, a four-way handshake occurs to authenticate the device. The pre-shared key and SSID are used to generate a session key during this process. A hacker can take advantage of some vulnerabilities in the four-way handshake to intercept the data and perform offline password attacks that could eventually crack weak passwords.

To address these vulnerabilities and support new technologies, the WPA3 standard is being implemented. WPA3 was introduced in 2018. Instead of using a pre-shared key, WPA3 uses the Simultaneous Authentication of Equals standard, or SAE standard.

SAE uses a 128-bit key and perfect forward secrecy to authenticate. Perfect forward secrecy is a cryptographic method that generates a new key for every transmission. This makes the handshake much more secure because if a hacker gets ahold of one message, they still aren't able to crack the keys.

It'll take a while for WPA3 to fully replace WPA2 since WPA2 is currently implemented in many network devices. Until then, we need to be aware of the differences in these standards and their encryption method vulnerabilities so that we can better protect our networks.

That'll wrap up this lesson. In this lesson, we looked at some things we can do to better secure out wireless networks. We first looked at some of the default configurations we should change, including login information, the SSID, MAC address filtering, and transmission power. Then we looked at the two main cryptographic protocols in use today, which are WPA2 and WPA3. WPA2 is still the most widely used protocol, but WPA3 has been released and will eventually fix many of the vulnerabilities inherent in WPA2.

8.6.2 Wireless Security Facts

Wireless networking uses radio frequencies to transmit data. This means anyone with a wireless receiver can capture data from an improperly secured network.

This lesson covers the following topics:

- Weak configurations
- Cryptographic protocols

Weak Configurations

Proper configuration of a wireless access point (WAP) is the first step in securing the network. The following table explains some important actions to take regarding WAP settings.

Security Configuration Action	Description

Change default login credentials	WAPs typically come configured with a default administrator username and password. Because the administrator username and password are used to configure WAP settings, it's important to reset the defaults. This prevents outsiders from guessing the default username and password and breaking into the system.
Change default service set identification (SSID) and broadcast	The SSID can be a maximum of 32 bytes in length. Since many manufacturers use a default SSID, it's important to change the SSID from the default. The SSID should be unique but should not contain identifiable information (address, last name, etc.). The SSID broadcast can also be disabled. This is known as SSID suppression or cloaking. A determined hacker can still easily discover hidden SSIDs. Disabling the SSID broadcast can cause connection issues for devices.
Enable MAC address filtering	Every network device has a unique media access control (MAC) address. By specifying the MAC addresses allowed to connect to the network, unauthorized MAC addresses can be prevented from connecting to the WAP. Configuring a MAC address filtering system is very time-consuming and demands upkeep. Attackers can still use tools to capture packets and retrieve valid MAC addresses. An attacker can spoof a wireless adapter's MAC address and circumvent the filter.
Update the firmware	Manufacturers release updates to the firmware on a regular basis to address known issues. It is important to regularly check for updates and apply them to prevent the system from being exposed to known bugs and security vulnerabilities. While it is extremely important to keep devices up-to-date, it's just as important to properly test new updates before pushing them out to the entire network. Proper testing will reduce the number of new bugs or problems on a live network that the update may have introduced.
Enable the WAP firewall	Most wireless APs come with a built-in firewall that connects the wireless network to a wired network. This should be enabled to help prevent unauthorized access to the network.
Wi-Fi signal strength	Data emanation is a significant security problem. By default, the radio signals used by a wireless network are broadcast omnidirectional and can travel quite a distance from the WAP. An attacker sitting outside the building may be able to connect to the wireless network if the signal is traveling outside. This can be limited by manipulating the WAP antenna placement. Some WAPs also allow the signal strength to be adjusted. These settings reduce the signal strength, so the signal stays inside the building.

Cryptographic Protocols

Enabling the proper cryptographic protocol is perhaps the most important way to secure a wireless network. For most users, Wi-Fi Protected Access (WPA) versions 2 or 3 will be the best option. The following table explains these two protocols:

Cryptographic Protocol	Description
Wi-Fi Protected Access 2 (WPA2)	WPA2 is the implementation name for wireless security that adheres to the 802.11i specifications. It was first introduced in 2004 and is still heavily used in today's networks. There are two versions of WPA2 available: WPA2-Personal is also known as WPA2-PSK (pre-shared key). This version uses a pre-shared key, or passphrase, to protect the network. WPA2-PSK: Uses Advanced Encryption Standard with Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (AES-CCMP)as the encryption algorithm to encrypt all data. AES-CCMP uses a 128-bit key and a 128-bit block size. Performs a 4-way handshake to authenticate the device when it connects to the access point. The pre-shared key and SSID are used to generate a session key during this process. The handshake does have some vulnerabilities that allow a hacker to intercept data and perform offline password attacks. WPA2-Enterprise uses a RADIUS server to authenticate users to the network.
Wi-Fi Protected Access 3 (WPA3)	To support the vulnerabilities inherent in the WPA2 handshake and to support newer technologies, WPA3 was implemented. First introduced in 2018, WPA3 implements the Simultaneous Authentication of Equals (SAE) standard instead of using the pre-shared key. SAE uses a 128-bit key and Perfect Forward Secrecy (PFS) to authenticate users. Perfect forward secrecy is a cryptography method that generates a new key for every transmission. This makes the handshake much more secure from hackers. If any portion of the handshake is intercepted, the key is still uncrackable.

8.6.3 Wireless Authentication and Access Methods (Lesson Video)

Transcript:

Wireless networks encrypt communications using a security protocol, which is typically WPA2 or WPA3. But in order to authenticate users and distribute authentication keys, other methods are generally used. In this lesson, we'll go over some of the different access methods we can implement to connect to a wireless network as well as the authentication protocols we often find in enterprise environments.

Depending on the environment, we can use different methods to connect users to a wireless network. Let's take a look at some of these methods.
The first and probably most common access method is to use a pre-shared key, or PSK. A PSK is simply a passphrase that we type in to connect to a network. If we don't want to type in the passphrase to connect each and every device, we can use Wi-Fi Protected Setup, or WPS, to simplify the process.
WPS only works with a network that uses a pre-shared key and WPA2. On the access point, there's usually a button that makes WPA2 start to search for devices in range. On the connecting device, there's probably a WPS button that automatically joins the device to the access point. You need to enter a unique eight-number access point pin if the connecting device doesn't have a button like this. Some devices and access points can also use Near Field Communication, or NFC, during the WPS process to connect to each other.

Our other option is to simply have an open network. This means that no authentication takes place, and anyone can connect to the network at will. This should only be used by public establishments that want to offer free wireless access. Many open networks implement a captive portal. Captive portals force a user to view and interact with them before accessing the network. Basically, you're initially able connect your device to the wireless network. But before you can access the internet, you're redirected to a captive portal page. You might be prompted to agree to the network's terms and conditions or maybe even asked to pay a fee before you can proceed. These portals are used a great deal today, and you've most likely used one before.

In an enterprise environment, using a PSK isn't very secure or efficient. Enterprise networks usually make use the 802.1x Protocol to authenticate users to the wireless network.

802.1x authentication is one of the most secure ways to enforce wireless network authentication. On a wired network, once a user is authenticated, the port they're connected to is activated. The port remains off if the user's activation fails. We can implement the 802.1x Protocol in a wireless network by enabling a virtual port when the user is authenticated. There are three components to an 802.1x setup. The first is the supplicant, or wireless client. Then we have the authenticator, which is responsible for handling the communications between the supplicant and the authentication serverâ€”the third component. Know that oftentimes, the authentication server is a RADIUS server. When you use a RADIUS server, the authenticator is also known as the network access server, or NAS.

RADIUS stands for Remote Authentication Dial-In Service. It was developed in 1991 and was originally used to authenticate users to remote networks over dial-up networks. We still use it today to remotely authenticate users. Obviously, though, we don't do this over dial-up connections anymore, which are more or less antiquated. RADIUS is known as a triple-A protocol, meaning that it provides authentication, authorization, and accounting management.

When using 802.1x with RADIUS, the client sends their credentials to the authenticator, or NAS. The NAS forwards the credentials to the RADIUS server to verify. If you use a Windows server, this is done using Active Directory. Then the server sends back the verification and user rights to the NAS, which forwards them back to the client. Then the client can access network resources. Using 802.1x authentication significantly increases your wireless network security.

To ensure that the authentication information being sent between these devices is secure, the Extensible Authentication Protocol, or EAP, is used. EAP isn't actually a specific protocol; it's merely a framework in which other protocols work. As a security administrator, you need to be familiar with several different EAP protocols. Let's start by looking at the Protected Extensible Authentication Protocol, or PEAP.

PEAP was created collaboratively by Cisco, Microsoft, and RSA Security. This protocol encapsulates the authentication communications within a TLS tunnel and exclusively uses a server-side certificate to authenticate Wi-Fi clients. This simplifies network implementation and administration.

Next, we have EAP Flexible Authentication via Secure Tunneling, or EAP-FAST. EAP-FAST was created by Cisco, and it works by making a TLS tunnel that doesn't require an authentication server certificate. Instead, this protocol uses a Protected Access Credential, or PAC, to authenticate users.

Then, we have EAP Transport Layer Security, or EAP-TLS. EAP-TLS is the original and probably most secure of the wireless EAP authentication protocols. Because it's the original, it's also the most widely supported.

The nice thing about EAP-TLS is that it requires client-side certificates in addition to server-side certificates. The certificate fully encrypts the authentication handshake between client and server and is then used in place of a password, making it practically impossible to crack. But because each client
requires an installed CA-signed PKI certificate, EAP-TLS is much more labor intensive and expensive to implement than other protocols.

The final protocol is EAP Tunneled Transport Layer Security, or EAP-TTLS. EAP-TTLS is essentially an updated version of EAP-TLS. The biggest difference is that we only need a certificate on the server. This greatly simplifies the implementation process since we don't need a certificate on each wireless device. EAP-TTLS has been supported natively since Windows 8. By using the right authentication protocol, you can greatly increase your wireless network's security and further protect it from attacks.

That'll wrap up this lesson. In this lesson, we looked at the different wireless network access methods we can implement, including pre-shared keys, server authentication, and open networks. We also went over how WPS and captive portals work. Then we saw the various authentication methods that can be used for wireless networks. First, we looked at RADIUS and how it utilizes 802.1x authentication. We finished up by looking at the different EAP protocols, which include PEAP, EAP-FAST, EAP-TLS, and EAP-TTLS.

8.6.4 Wireless Authentication and Access Methods Facts

This lesson covers the following topics:

Access methods

Port security

MAC filtering and MAC limiting

Access Methods

Choose an access method for a wireless network based on the use of the network. The following table describes access methods:

Access Method	Description
Pre-shared key (PSK)	A pre-shared key is a passphrase used to access the wireless network. This is probably the most commonly used access method.
Wi-Fi Protected Setup (WPS)	Wi-Fi Protected Setup works only on a network that uses a PSK and WPA2. WPS allows a device to securely connect to a wireless network without typing in the PSK. To do this, you: Push a button on the access point that causes the access point to search for devices in range. Push the WPS button on the device to automatically join it to the access point. If there is no button, enter the eight-digit pin unique to the access point. Some devices and access points can also use Near Field Communication (NFC) during the WPS process to connect to each other.
Open network	An open network has no authentication at all and allows anyone to connect to the network. This access method should be used only in public places that want to offer free wireless access.
Captive portal	Many open networks implement a captive portal. Captive portals force a user to view and interact with them before accessing a network. A hotel network is a good example of captive portal use. When using a captive portal: The user connects to the wireless network but is redirected to a captive portal page before internet access is granted. The user might be prompted to agree to the terms and conditions of using the network or even asked to pay a fee before being granted access.

To ensure the authentication information being sent is secure, the Extensible Authentication Protocol (EAP) is used. EAP is a framework in which other protocols work. The following table explains EAP and the protocols:

Protocol	Description

Extensible Authentication Protocol (EAP)	EAP is a set of interface standards that allows various authentication methods to be used: • EAP supports multiple authentication methods (smart cards, biometrics, and digital certificates). • Using EAP, the client and server negotiate the characteristics of authentication.
Protected Extensible Authentication Protocol (PEAP)	PEAP provides authentication in an SSL/TLS tunnel with a single certificate on the server. PEAP: • Creates a secure communication channel for transmitting certificate or login credentials. • Enables mutual authentication by requiring the server to prove its identity to the client. • Was a collaborative effort between Cisco, Microsoft, and RSA.
EAP Flexible Authentication via Secure Tunneling (EAP-FAST)	EAP-FAST uses a Protected Access Credential (PAC) to authenticate users. EAP-FAST: • Establishes a TLS tunnel in which client authentication credentials are transmitted. • Is susceptible to attackers who intercept the Protected Access Credential (PAC) and use it to compromise user credentials. This vulnerability is mitigated by manual PAC provisioning or by using server certificates. • Was created by Cisco.
EAP Transport Layer Security (EAP-TLS)	EAP-TLS uses Transport Layer Security (TLS) and is considered one of the most secure EAP standards available. EAP-TLS: • Is widely supported by almost all manufacturers of wireless LAN hardware and software. • Requires signed client-side and server-side certificate authority (CA) PKI certificates. • Is labor-intensive and expensive to implement.
EAP Tunneled Transport Layer Security (EAP-TTLS)	EAP-TTLS also uses a CA signed certificate. EAP-TTLS: • Is an updated version of EAP-TLS. • Requires only one CA signed certificate on the server, simplifying the implementation process.

Port Security

Each wall and switch port represents an opportunity for a threat actor to attach a device to the network. A threat actor who can operate a host with physical access to a network segment can launch a variety of attacks.

Access to the physical switch ports and switch hardware should be restricted to authorized staff. To accomplish this, place the switch appliances in secure server rooms and lockable hardware cabinets. To prevent the attachment of unauthorized client devices at unsecured wall ports, the switch port that the wall port cabling connects to can be administratively disabled, or the patch cable can be physically removed from the switch port. Completely disabling ports in this way can introduce a lot of administrative overhead and allow room for error. Also, it doesn't provide complete protection, as an attacker could unplug a device from an enabled port and connect their own machine. Consequently, more sophisticated methods of ensuring port security have been developed.

MAC Filtering and MAC Limiting

The network adapter in each host computer is identified by a MAC address. Configuring MAC filtering means a switch port only permits certain MAC addresses to connect. This can be done by creating a list of valid MAC addresses or by specifying a limit to the number of permitted addresses. For example, if port security is enabled with a maximum of two MAC addresses, the switch will record the first two MACs to connect to that port. The switch then drops any traffic from machines with different MAC addresses that try to connect.

```
NYCORE1>
NYCORE1#
*Mar  1 00:02:27.991: %SYS-5-CONFIG_I: Configured from console by console
*Mar  1 00:02:46.287: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to up
NYCORE1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
NYCORE1(config)#ip arp inspection vlan 1,999
NYCORE1(config)#
*Mar  1 00:07:20.561: %SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa1/0/23, vlan 1.([0023.049
0.0000/192.168.16.21/00:07:20 UTC Mon Mar 1 1993])
```

Configuring ARP inspection on a Cisco switch. (Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)

Restricting access by MAC address is difficult to manage and still prone to spoofing. Better security is obtained by forcing computers and users to authenticate before full network access is granted. The IEEE 802.1X Port-based Network Access Control (P NAC) standard allows a switch to require authentication when a host connects to one of its ports. 802.1X uses authentication, authorization, and accounting (AAA) architecture:

Supplicant - is the device requesting access, such as a user's PC or laptop.

Authenticator - is the switching device (or any type of network access appliance). This does not validate authentication request s directly but acts as a conduit for authentication data.

Authentication server - the server that holds or can contact a directory of network objects and validate authentication requests, issue authorizations, and perform accounting of security events.

The 802.1X standard is implemented by two protocols:

Extensible Authentication Protocol (EAP) - provides a framework for deploying multiple types of authentication methods. It is often used with digital certificates to establish a trust relationship and create a secure tunnel to transmit the user credential or to perform smart-card authentication without a password.

Remote Authentication Dial-In User Service (RADIUS) - allows the authenticator and authentication server to communicate authentication and authorization decisions. The authenticator is a RADIUS client; the authentication server is a RADIUS server.

When a host connects to an 802.1X-enabled switch port, the switch opens the port for the EAP over LAN (EAPoL) protocol only. The switch port only allows full data access when the host has been authenticated. The switch receives an EAP packet with the supplicant's credentials. These are encrypted and cannot be read by the switch. The switch uses the RADIUS protocol to send the EAP packet to the authentication server. The authentication server can access the directory of user accounts and can validate the credential. If authentication is successful, it informs the switch that full network access can be granted.

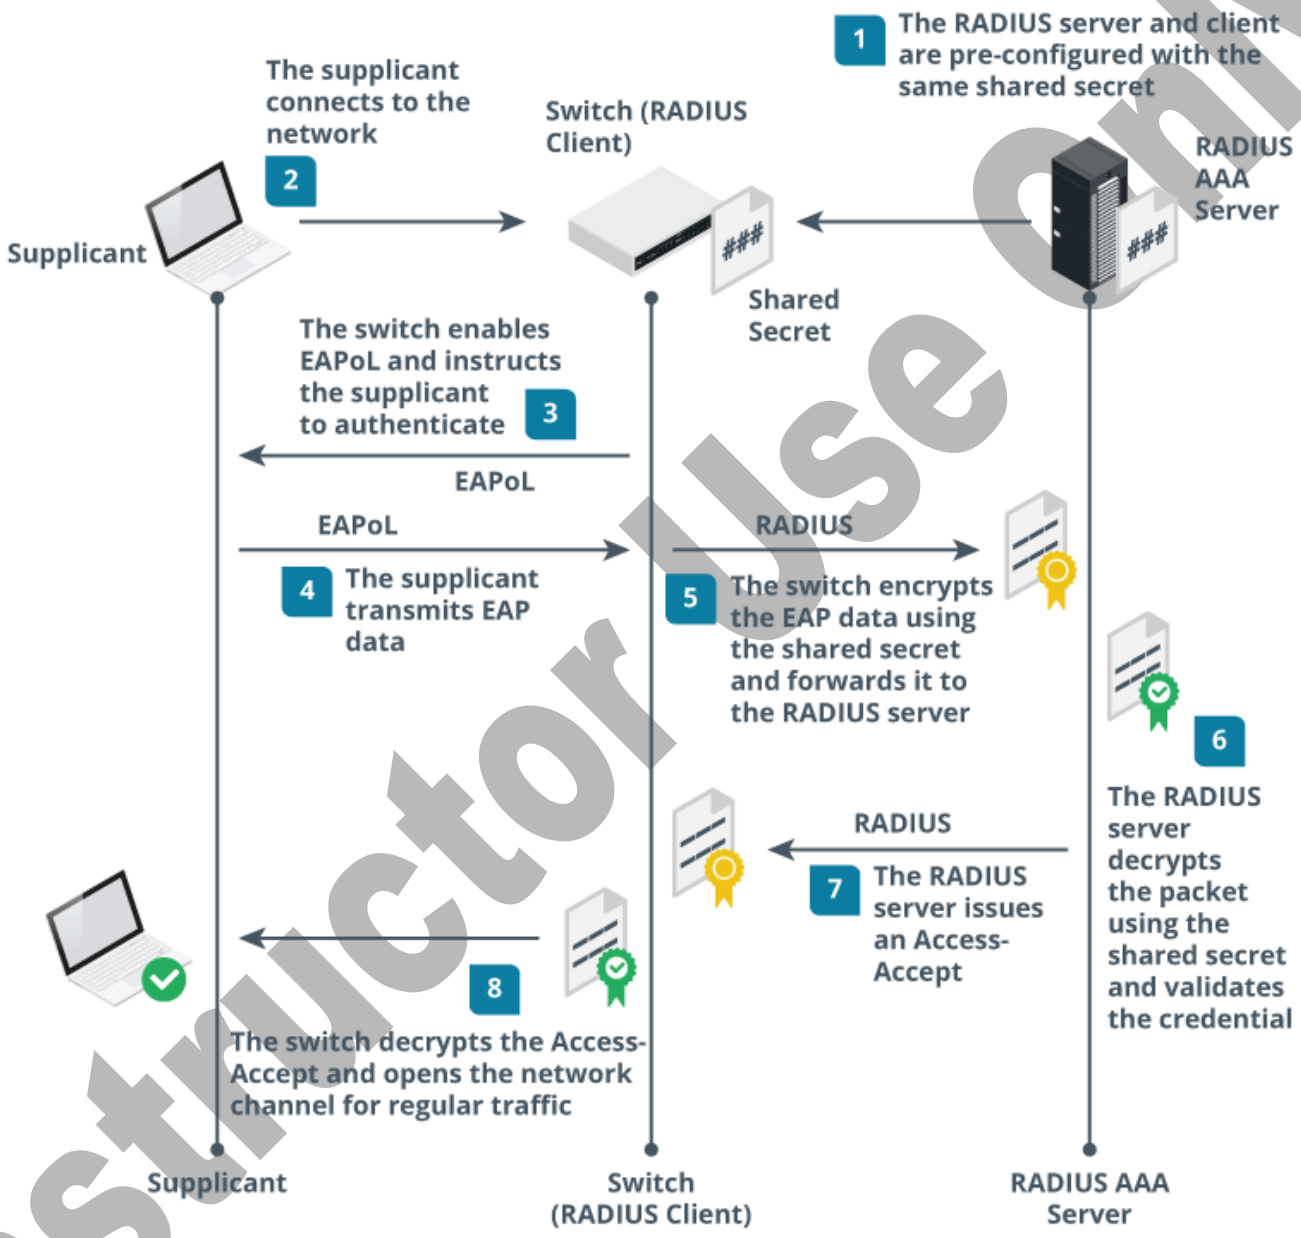

IEEE 802.1X Port-based Network Access Control with RADIUS and EAP authentication. (Images © 123RF.com.)

Some hosts are so security-critical that it is unsafe to connect them to any type of network. One example is the root certification authority in PKI. Another example is a host used to analyze malware execution. A host not physically connected to any network is said to be air-gapped .

It is also possible to configure an air-gapped network. This means that hosts within the air-gapped network can communicate, but there is no cabled or wireless connection to any other network. Military bases, government sites, and industrial facilities use air-gapped networks.

Physically isolating a host or group of hosts improves security but also incurs significant management challenges. Device administration has to be performed at a local terminal. Any updates or installs have to be performed using USB or optical media. This media is a potential attack vector and must be scanned before allowing its use on an air-gapped host.

8.6.5 Hardening a Wireless Access Point (Demo Video)

Transcript:

In this demonstration, we're going to go through some ofÂ the steps you can take to harden a wireless network, andÂ we're going to use a wireless controller.

This means we make changes on the controllerÂ itself, and those changes propagate out to eachÂ of the access points connected to it.Â The concepts in this demo transfer toÂ other types of wireless access points, such asÂ all-in-one small office/home office devices.
To start, we need to log in to our controller.Â I have a web browser open, and I've entered the controller's IP address. I'll enterÂ the username, â€˜admin', and password, and we're logged in.
The first thing we want to do is change theÂ default username and password for this device.Â This is the first thing you should do withÂ any networking device you're configuring.Â It's very important that you change those defaults.
We do this on the Administer tab.Â Notice we can also choose to use anÂ authorization server if we want to, but justÂ change the default username to something else.Â We'll enter the old password and then the new password.Â Click Apply. The Admin session timeout setting is also on this page.
This setting defines how much time must passÂ before the user is automatically logged out.
This is a nice thing to have, as it keeps users from logging in to the device andÂ then walking away from it and leaving it vulnerable.Â Let's change this to 10 minutes.Â Click Apply.
Now we're asked to log out and start aÂ new session with the new timeout setting.Â Click OK, and let's log in withÂ our new username and password.
Let's go to the Configure tab.Â There are a lot of settings and options on this tab.Â On the System tab, we can specify the IPÂ address and subnet mask to use on the network.
It's good to change this from a typical IPÂ address range, as it adds another layer ofÂ security, but we'll leave it how it is for now.Â Let's go to the Access Points page.
On this page, we can see all the accessÂ points connected to the wireless controller.
Right now, we only have one AP, or access point, connected.Â If we wanted to, we could clickÂ Edit and change the configuration.
Notice that all these options requireÂ us to override the group configuration.Â Remember, this is a wireless controller.Â As such, it controls the configurationÂ options for each access point in the network.
If we want to change the settings on aÂ specific access point (for example, change theÂ channel it's using), we can do that here.Â For now, let's click Cancel.
Under Access Point Groups, any changes to thisÂ configuration will be applied to all the APs in the network.Â This makes managing networks with multiple APs very easy.Â All the settings look good, so let's just click Cancel.
On the WLANs page, we can create wirelessÂ LANs--that is, we can create multipleÂ wireless networks, each specified by an SSID.Â For example, we currently have threeÂ different WLANs configured on this controller.Â Let's edit the DevRuckusAP network.
Here, we can see the configuration of this networkÂ and that it's defined as a standard network.Â We can also configure the authentication options.Â Right now, the authenticate type is open.
We could configure this to connect to a RADIUSÂ server, specify that it use MAC addressÂ authentication, or use a combination of the two.Â We can also configure the encryption method.
We're using WPA2 with the AES algorithm.Â You can see we have a very weak passphrase configured,Â which is fine for a demo, but make sure you use a complex passphrase with this encryption.
Down here, we can also configureÂ the WLAN to use a captive portal.Â The Wireless Client Isolation option allows us toÂ configure the network to isolate each connected client.

There are also a lot of advanced options we won't getÂ into, but remember there's a lot of customization thatÂ can happen with these types of wireless controllers thatÂ you just can't find in smaller, consumer-grade devices.Â Leave everything as it is and click Cancel.
Let's look at the configuration of the Guest network.Â Notice the type is defined as Guest, which causes our guestÂ access policies and access control settings to be applied.
We also need to note that we're using openÂ authentication with no encryption.Â This is an open, unsecure guest network.Â That's why we're isolatingÂ wireless client traffic from each other.Â It's good idea to have thisÂ option enabled on guest networks.Â Click Cancel, and let's go to the Access Control page.
On this page, we can configureÂ how client access is controlled.Â For example, under L2-L7 Access Control, we can define MACÂ address control lists that can be used with a specific WLAN.
To define MAC address control lists, add all the MACÂ addresses you want to give network access to and then enableÂ MAC address filtering on the specific WLAN to use that list.Â We can also do this by IP address.
Under Device Access Policy, we can permit or denyÂ clients based on their device specifications.Â For example, we could deny all clientsÂ using a particular operating system.
Application Recognition and Filtering is another useful feature.Â It lets us block traffic based onÂ port, IP address, protocol, and other characteristics.
We configure this by looking at our access points'Â logs and identifying the different types of traffic.Â We would then take that information andÂ decide if we want to block anything.
Let's go back up to the deviceÂ access policy and define a rule.Â We're going to create a rule thatÂ blocks traffic from video game consoles.Â We'll name the rule and add a description.Â Click Create New.
Notice, under OS/Type, that we have an option for gaming.Â This identifies traffic generated by videoÂ game consoles, such as an Xbox or PlayStation console.Â We want to deny this traffic, so we'llÂ select Deny from the rule type dropdown.Â Click Save > OK.
There's a number of other featuresÂ on this controller that you can use.Â We can add roles and policies that affectÂ users, and we can add users to the device.Â We can also configure the controller toÂ notify us of specific events via email.
On the Services page, we can configure self-healingÂ and background scanning, which are both very usefulÂ feature, but they're outside the scope of this video.Â We can also configure wireless intrusion detection.
Wireless intrusion detection and preventionÂ system (WIPS) is, actually, a very useful feature.Â It helps protect against wireless attacks such asÂ denial-of-service (DoS) attacks, rogue access points,Â evil twin attacks, and rogue DHCP servers.
For example, we can protect againstÂ repeated authentication attacks right here.Â We can even specify how long a client will be blocked.Â This helps protect our system from DoSÂ attacks and brute force attacks. We can also configure the controllerÂ to automatically flag APs trying to spoof anÂ SSID or MAC address as a malicious device.Â This is a great tool for protectingÂ against evil twin attacks.
That's it for this demonstration.Â In this demo, we discussed some of theÂ things you can do to harden a wirelessÂ network using a wireless network controller.

8.6.6 Harden a Wireless Network (Simulation)

Scenario

You are a network technician for a small corporate network. You need to increase the security of your wireless network. Your new wireless controller provides several security features that you want to implement.

Access the Wireless Controller console through Chrome on **http://192.168.0.6** with the username **admin** and the password **password** . The username and password are case-sensitive.

In this lab, your task is to:

Change the admin username and password for the Zone Director controller to the following:

Admin Name: **WxAdmin**

Password: **ZDAdminsOnly!$** (O is the capital letter O)

Set up MAC address filtering (L2 Access Control) to create an allow list called **Allowed Devices** that includes the following wireless devices:

00:18:DE:01:34:67

00:18:DE:22:55:99

00:02:2D:23:56:89

00:02:2D:44:66:88

Implement a device access policy called **NoGames** that blocks gaming consoles from the wireless network.

Explanation

Complete this lab as follows:

Access the Ruckus zone controller.

From the taskbar, select **Google Chrome** .

In the URL field, enter **192.168.0.6** and press **Enter** .

Maximize the window for better viewing.

Log in to the wireless controller console.

In the Admin field, enter **admin** (case sensitive).

In the Password field, enter **password** as the password.

Select **Login** .

Change the admin username and password for the Zone Director controller.

From the top, select the **Administer** tab.

Make sure **Authenticate using the admin name and password** is selected.

In the Admin Name field, enter **WxAdmin** .

In the Current Password field, enter **password** .

In the New Password field, enter **ZDAdminsOnly!$** .

In the Confirm New Password field, enter **ZDAdminsOnly!$** .

On the right, select **Apply** .

Enable MAC address filtering.

From the top, select the **Configure** tab.

From the left menu, select **Access Control** .

Expand **L2-L7 Access Control** .

Under *L2/MAC address Access Control* , select **Create New** .

In the Name field, enter **Allowed Devices** .

Under Restriction, make sure **Only allow all stations listed below** is selected.

Enter a ***MAC address*** .

Select **Create New** .

Repeat step 4g–4h for each MAC address you would like to add to the ACL.

Select **OK** .

Configure access controls.

Under Access Control, expand **Device Access Policy** .

Select **Create New** .

In the Name field, enter **NoGames** .

Select **Create New** .

In the Description field, enter **Games** .

Using the OS/Type drop-down list, select **Gaming** .

In the Type field, select **Deny** .

Under Uplink, make sure **Disabled** is selected.

Under Downlink, make sure **Disabled** is selected.

Select **Save** .

Select **OK** .

8.6.7 Configure WIPS (Simulation)

Scenario

You are a network technician for a small corporate network. You would like to enable Wireless Intrusion Prevention on the wireless controller. You are already logged in as WxAdmin.

Access the Wireless Controller console through Chrome on **http://192.168.0.6** .

In this lab, your task is to:

Configure the wireless controller to protect against denial-of-service (DOS) attacks as follows:

Protect against excessive wireless requests.

Block clients with repeated authentication failures for two minutes (120 seconds).

Configure Intrusion Detection and Prevention as follows:

Report all rogue devices regardless of type.

Protect the network from rogue access points.

Enable **Rogue DHCP Server Detection** .

Explanation

Complete this lab as follows:

Access the Ruckus zone controller.

From the taskbar, select **Google Chrome** .

In the URL field, enter **192.168.0.6** and press **Enter** .

Maximize the window for better viewing.

Configure Denial of Service protection.

Select the **Configure** tab.

From the left menu, select **WIPS** .

From the right, under *Denial of Services (DoS)* , select **Protect my wireless network against excessive wireless requests** .

Select **Temporarily block wireless clients with repeated authentication failures** .

Enter **120** seconds.

From the right of this section, select **Apply** .

Configure Intrusion Detection and Prevention:

Under *Intrusion Detection and Prevention* , select **Enable report rogue devices** .

Select **Report all rogue devices** .

Select **Protect the network from malicious rogue access points** .

From the right of this section, select **Apply** .

Configure Rogue DHCP Server Detection.

Under *Rogue DHCP Server Detection* , select **Enable rogue DHCP server detection** .

From the right of this section, select **Apply** .

8.6.8 Implement Secure Wireless Infrastructure

8.6.9 Practice Questions (Section Quiz)

q_wl_security_psk_secp8

You want to connect a laptop computer running Windows to a wireless network.

The wireless network uses multiple access points and WPA2-Personal. You want to use the strongest authentication and encryption possible. SSID broadcast has been disabled.

What should you do?

Answers:

***Configure the connection with a pre-shared key and AES encryption.**

Configure the connection to use 802.1x authentication and AES encryption.

Configure the connection with a pre-shared key and TKIP encryption.

Configure the connection to use 802.1x authentication and TKIP encryption.

Explanation:

To connect to a wireless network using WPA2-Personal, you need to use a pre-shared key for authentication. Advanced Encryption Standard (AES) encryption is supported by WPA2 and is the strongest encryption method.

WPA and WPA2 designations that include Personal or PSK use a pre-shared key for authentication.

Methods that include Enterprise use a RADIUS server for authentication and 802.1x authentication with usernames and passwords.

q_wl_security_update_firmware_secp8

A company's wireless network has been experiencing intermittent connectivity issues and slower than usual data transfer speeds.

The network administrator recently updated the firmware on the wireless access point (WAP) as part of a routine maintenance procedure. The update was applied during a scheduled downtime and the network was functional when the downtime ended.

However, the issues started appearing the next day. The administrator has checked the WAP settings and everything seems to be in order.

What could be the MOST likely cause of these issues?

Answers:

The WAP's firewall was accidentally disabled during the update.

***The network administrator failed to properly test new updates before pushing them out to the network.**

The MAC address filtering system was not configured correctly.

The Wi-Fi signal strength was set too high, causing interference.

Explanation:

The correct answer is most likely that the network administrator failed to properly test new updates before pushing them out to the network. Not properly testing new updates before pushing them out to the network could lead to unforeseen issues, such as the ones being experienced. The update could have introduced bugs or incompatibilities that were not apparent immediately after the update was applied.

The WAP's firewall being disabled could potentially allow unauthorized access to the network, but it would not typically cause intermittent connectivity issues or slower data transfer speeds.

Incorrect configuration of the MAC address filtering system could prevent certain devices from connecting to the network, but it would not typically cause intermittent connectivity or slower data transfer speeds for devices that are able to connect.

While a high Wi-Fi signal strength could potentially cause interference, it would not typically cause the specific issues being experienced unless there were other networks nearby on the same channel. Additionally, this would not explain why the issues started appearing after the firmware update.

q_wl_security_wpa2_01_secp8

Which of the following features is supplied by WPA2 on a wireless network?

Answers:

***Encryption**

Client-connection refusal based on MAC address

Traffic filtering based on packet characteristics

Network identification

Centralized access point for clients

Explanation:

Wi-Fi Protected Access (WPA) provides encryption and user authentication for wireless networks.

MAC address filtering allows or rejects client connections based on the hardware address.

The SSID is the network name or identifier.

A wireless access point (called an AP or WAP) is the central connection point for wireless clients.

A firewall allows or rejects packets based on packet characteristics (such as address, port, or protocol type).

q_wl_security_wpa2_02_secp8

You need to secure your wireless network.

Which security protocol would be the BEST choice?

Answers:

WEP

WPA

EFS

***WPA2**

802.11n

Explanation:

WEP, WPA, and WPA2 are all security protocols for wireless networks. Each security protocol protects the wireless data through the use of association keys and encryption protocols. However, WPA2 provides the best wireless security.

802.11n is a wireless standard with specific parameters for wireless data transmission.

The Encrypting File System (EFS) is a method for encrypting individual files within Windows.

q_wl_security_wpa2_03_secp8

You need to add security for your wireless network, and you would like to use the most secure method.

Which method should you implement?

Answers:

WEP

WPA

***WPA2**

Kerberos

Explanation:

Wi-Fi Protected Access 2 (WPA2) is currently the most secure wireless security specification. WPA2 includes specifications for both encryption and authentication.

WPA was an earlier implementation of security specified by the 802.11i committee. WEP was the original security method for wireless networks. WPA is more secure than WEP but less secure than WPA2.

Kerberos is an authentication method, not a wireless security method.

q_wl_auth_access_captive_secp8

The owner of a hotel has contracted with you to implement a wireless network to provide internet access for guests.

The owner has asked that you implement security controls so that only paying guests are allowed to use the wireless network. She wants guests to be presented with a login page when they initially connect to the wireless network. After entering a code provided by the concierge at check-in, guests should then be allowed full access to the internet. If a user does not provide the correct code, he or she should not be allowed to access the internet.

What should you do?

Answers:

***Implement a captive portal**

Implement MAC address filtering

Implement 802.1x authentication using a RADIUS server

Implement pre-shared key authentication

Explanation:

A captive portal would be the best choice in this scenario. A captive portal requires wireless network users to abide by certain conditions before they are allowed access to the wireless network. For example, the captive portal could require them to:

Agree to an acceptable use policy

Provide a PIN or password

Pay for access to the wireless network

View information or advertisements about the organization providing the wireless network (such as an airport or hotel)

When a wireless device initially connects to the wireless network, all traffic to or from that device is blocked until the user opens a browser and accesses the captive portal webpage. After the user provides the appropriate code, traffic is unblocked, and the host can access the network normally.

MAC address filtering and 802.1x authentication would work from a technical standpoint, but these would be completely unmanageable in a hotel scenario where guests come and go every day. Using a pre-shared key would require a degree of

technical expertise on the part of the hotel guests. It could also become problematic if the key were to be leaked, allowing non-guests to use the wireless network.

q_wl_auth_access_eap_802_01_secp8

An educational institution plans to upgrade its Wi-Fi network to improve security and control who can access the network while maintaining ease of connection for a large number of diverse users.

What is the BEST practice for this institution to improve its Wi-Fi network's security and manageability?

Answers:

***Implement the 802.1X standard using Extensible Authentication Protocol (EAP) for user authentication.**

Use a pre-shared key (PSK) for all devices.

Implement an open network and monitor network activity for unusual behavior.

Apply media access control (MAC) address filtering and maintain an authorized list of MAC addresses.

Explanation:

The network security administrator implements the 802.1X standard using EAP as best practice for enterprise Wi-Fi security. This allows for individual user authentication, which means the network provides access control at the user level.

The administrator might find using a PSK for all devices convenient, but it doesn't offer individual user accountability.

The administrator exposes the network to significant risk when they set up an open network and monitor activity. Unusual behavior is only detected after a security incident since any device can connect to the network.

The administrator applies MAC address filtering to prevent unauthorized devices from connecting, but managing this can be time-consuming with many users, and it is susceptible to MAC spoofing.

q_wl_auth_access_eap_802_02_secp8

A multinational corporation with numerous remote workers considers methods to enhance its network security. It wants to ensure that only authorized employees can connect their devices to the company network, regardless of location.

Which solution should the corporation implement to securely authenticate remote workers and maintain robust port security?

Answers:

***Employ 802.1X with Extensible Authentication Protocol (EAP) for robust port security and user authentication.**

Deploy a network that is open to all and regularly audit log files for suspicious activities.

Implement a media access control (MAC) address authorized list that only allows approved devices to connect to the network.

Use static Internal Protocol (IP) addresses for all employees and deny network access to devices with unrecognized IPs.

Explanation:

The network security administrator employs 802.1X with Extensible Authentication Protocol (EAP) for secure and robust user-level authentication. This works with a range of EAP methods depending on the corporation's specific needs and provides robust port security.

If the administrator sets up an open network and monitors and audits it for suspicious activities, it shows a lack of strong access control measures.

Maintaining a MAC address authorized list can be challenging for the administrator, especially in a multinational corporation with many employees. This approach also doesn't provide user-level access control.

The administrator finds using static IP addresses ineffective for remote workers who often need to have static IPs at their locations. This approach also doesn't provide any form of user authentication.

q_wl_auth_access_eap_radius_secp8

The security team at a multinational cloud services company is working on their port security. They implemented basic Media Access Control (MAC) address filtering on all switch ports, but they have concerns about the risk of MAC spoofing and the management overhead of maintaining a list of valid MAC addresses.

To address these concerns, they now require strong authentication before a user can obtain full network access.

Which of the following measures should the team implement next?

Answers:

***Implement EAP and RADIUS.**

Implement physical isolation for critical servers.

Disable all unused switch ports.

Segregate into multiple security zones.

Explanation:

Implementing IEEE 802.1X standard with extensible authentication protocol (EAP) and Remote Authentication Dial-In User Service (RADIUS) mitigates the risk of MAC spoofing and reduces management overhead of maintaining a list of valid MAC addresses. The switch requires authentication when a host connects to one of its ports, providing secure port access.

Physical isolation can provide strong security for critical servers, but it does not address the need for authenticated access at the port level.

While disabling unused ports can limit unauthorized physical access, it does not solve the problem of MAC spoofing or the overhead of maintaining MAC filtering.

Creating additional security zones can help limit the scope of potential network breaches but does not directly address port security.

q_wl_auth_access_eap_secp8

Which remote access authentication protocol allows for the use of smart cards for authentication?

Answers:

***EAP**

CHAP

PAP

PPP

SLIP

Explanation:

Extensible Authentication Protocol (EAP) is a set of interface standards that allows you to use various authentication methods, including smart cards, biometrics, and digital certificates.

Password Authentication Protocol (PAP) transmits login credentials in cleartext.

Challenge Handshake Authentication Protocol (CHAP) protects login credentials using a hash and allows for periodic re-authentication.

Point-to-Point Protocol (PPP) and Serial Line Interface Protocol (SLIP) are not remote access authentication protocols. They are used to establish the connection, not provide authentication.

q_wl_auth_access_eap_tls_01_secp8

Which EAP implementation is MOST secure?

Answers:

***EAP-TLS**

EAP-MD5

EAP-FAST

LEAP

Explanation:

EAP-TLS uses Transport Layer Security (TLS) and is considered one of the most secure EAP standards available. A compromised password is not enough to break into EAP-TLS enabled systems because the attacker must also have the client's private key.

EAP-MD5 offers minimal security and is susceptible to dictionary attacks and man-in-the-middle attacks. Lightweight Extensible Authentication Protocol (LEAP) does a poor job of protecting user authentication credentials and is also susceptible to dictionary attacks.

EAP-FAST is a replacement for LEAP that uses a protected access credential (PAC) to establish a TLS tunnel in which client authentication credentials are transmitted.

While more secure than EAP-MD5 and LEAP, EAP-FAST can still be compromised if the attacker intercepts the PAC.

q_wl_auth_access_eap_tls_02_secp8

The cybersecurity team at an international data center operator is enhancing its port security strategy. They have implemented rudimentary media access control (MAC) address filtering on all switch ports.

However, concerns about MAC spoofing and the administrative burden of maintaining a list of valid MAC addresses have surfaced. The team decided that robust authentication needs to occur before a user can gain comprehensive network access.

What strategy should the cybersecurity team adopt next?

Answers:

***Deploy EAP over TLS.**

Enforce physical isolation for servers.

Enable auxiliary switch ports.

Divide into secure segments.

Explanation:

Deploying the IEEE 802.1X standard with EAP over TLS effectively addresses concerns about MAC spoofing and the administrative overhead of maintaining a list of valid MAC addresses. It also switches mandates authentication when a host connects to one of its ports, providing secure port access.

Physical isolation provides robust security for backup servers but does not address the requirement for authenticated access at the port level.

Enabling auxiliary ports can open up more points for unauthorized access and does not address the problem of MAC spoofing or the overhead of maintaining MAC filtering.

Adding more secure segments can help limit the reach of potential network breaches but does not directly address port security.

q_wl_auth_access_mac_secp8

Which of the following do switches and wireless access points use to control access through a device?

Answers:

***MAC address filtering**

IP address filtering

Port number filtering

Session filtering

Explanation:

Both switches and wireless access points are Layer 2 devices, meaning they use the MAC address to make forwarding decisions. Both devices typically include some form of security that restricts access based on the MAC address.

Routers and firewalls operate at Layer 3 and can use the IP address or port number for filtering decisions.

A circuit-level gateway is a firewall that can make forwarding decisions based on the session information.

q_wl_auth_access_peap_01_secp8

The IT team of a large multinational corporation is working to improve the security of their remote access services. They plan to implement Remote Authentication Dial-In User Service (RADIUS) to enhance the authentication process for remote users. RADIUS provides a centralized authentication and authorization mechanism for users connecting from various locations.

The IT team evaluated different authentication protocols alongside RADIUS to ensure a strong and secure remote access solution.

Which choice of authentication protocols would be MOST appropriate to complement RADIUS for the company's remote access solution?

Answers:

***Protected Extensible Authentication Protocol (PEAP)**

Password Authentication Protocol (PAP)

Wired Equivalent Privacy (WEP)

Address Resolution Protocol (ARP)

Explanation:

Protected Extensible Authentication Protocol (PEAP) gets widely used with RADIUS for remote access authentication. It creates a secure transport layer security (TLS) tunnel for transmitting user credentials and provides a robust authentication method for remote users.

Password Authentication Protocol (PAP) is weak as it transmits passwords in plaintext, making it vulnerable to interception and unauthorized access.

Wired Equivalent Privacy (WEP) is an outdated and weak security protocol easily exploited, making it unsuitable for a secure remote access solution.

Address Resolution Protocol (ARP) is not an authentication protocol. ARP relates to mapping internet protocol (IP) addresses to media access control (MAC) addresses in local networks and is irrelevant to remote access security.

q_wl_auth_access_peap_02_secp8

The IT team of a medium-sized governmental agency took significant steps to bolster the security of their remote access services. Determined to fortify the authentication process for remote users, the team decided to implement Remote Authentication Dial-In User Service (RADIUS).

However, relying solely on RADIUS may not be enough to achieve their rock-solid remote access solution goal. To ensure the utmost protection, the team evaluated different authentication protocols that can complement and work seamlessly with RADIUS.

What authentication protocols should the team use to complement RADIUS for the company's remote access solution?

Answers:

***Protected Extensible Authentication Protocol (PEAP)**

Password Authentication Protocol (PAP)

Wired Equivalent Privacy (WEP)

Address Resolution Protocol (ARP)

Explanation:

Protected Extensible Authentication Protocol (PEAP) gets widely used with RADIUS for remote access authentication. It creates a secure transport layer security (TLS) tunnel for transmitting user credentials and provides a robust authentication method for remote users.

Password Authentication Protocol (PAP) is weak as it transmits passwords in plaintext, making it vulnerable to interception and unauthorized access.

Wired Equivalent Privacy (WEP) is an outdated and weak security protocol easily exploited, making it unsuitable for a secure remote access solution.

Address Resolution Protocol (ARP) is not an authentication protocol. ARP relates to mapping internet protocol (IP) addresses to media access control (MAC) addresses in local networks and is irrelevant to remote access security.

q_wl_auth_access_port_security_secp8

Which of the following methods can be used to ensure port security by completely disabling access to a network?

Answers:

Implementing IEEE 802.1X Port-based Network Access Control

Using a captive portal

***Physically removing the patch cable from the switch port**

MAC filtering

Explanation:

Physically removing the patch cable from the switch port is the correct answer. This is a method of port security where access to the physical switch ports and switch hardware is restricted to authorized staff. The switch port that the wall port cabling connects to can be administratively disabled, or the patch cable can be physically removed from the switch port. This method completely disables access to the network.

Implementing IEEE 802.1X Port-based Network Access Control is a method of port security where a switch requires authentication when a host connects to one of its ports. It uses authentication, authorization, and accounting (AAA) architecture to control access. However, this method does not completely disable access to the network.

Using a captive portal is a method used in open networks to force users to agree to terms and conditions or to pay a fee before being granted access. It does not disable access to the network, but rather controls it.

MAC filtering is a method of port security where a switch port only permits certain MAC addresses to connect. This can be done by creating a list of valid MAC addresses or by specifying a limit to the number of permitted addresses. However, this method does not completely disable access to the network.

q_wl_auth_access_radius_01_secp8

You want to implement 802.1x authentication on your wireless network.

Which of the following is required?

Answers:

WPA

WPA2

***RADIUS**

TKIP

Explanation:

802.1x authentication uses usernames and passwords, certificates, or devices such as smart cards to authenticate wireless clients. 802.1x authentication requires the following components:

- A RADIUS server to centralize user account and authentication information. A centralized database for user authentication is required to allow wireless clients to roam between cells but authenticate using the same account information.
- A public key infrastructure (PKI) for issuing certificates. At a minimum, the RADIUS server must have a server certificate. To support mutual authentication, each client must also have a certificate.

You can use 802.1x authentication with both WPA and WPA2 and even with WEP on some devices and operating systems. Temporary Key Integrity Protocol (TKIP) is an encryption method used with WPA.

q_wl_auth_access_radius_02_secp8

You want to implement 802.1x authentication on your wireless network.

Where would you configure passwords that are used for authentication?

Answers:

***On a RADIUS server**

On the wireless access point

On the wireless access point and on each wireless device

On a certificate authority (CA)

Explanation:

802.1x authentication uses usernames and passwords, certificates, or devices such as smart cards to authenticate wireless clients. Authentication requests received by the wireless access point are passed to a RADIUS server, which validates the login credentials (such as the username and password).

If you are using pre-shared keys for authentication, configure the same key on the wireless access point and on each wireless device.

A CA is required to issue a certificate to the RADIUS server.

The certificate proves the identity of the RADIUS server and can also be used to issue certificates to individual clients.

q_wl_auth_access_radius_03_secp8

You are the wireless network administrator for your organization. As the size of the organization has grown, you've decide to upgrade your wireless network to use 802.1x authentication instead of pre-shared keys.

To do this, you need to configure a RADIUS server and RADIUS clients. You want the server and the clients to mutually authenticate with each other.

What should you do? (Select two. Each response is a part of the complete solution.)

Answers:

***Configure the RADIUS server with a server certificate.**

***Configure all wireless access points with client certificates.**

Configure all RADIUS clients with a pre-shared key.

Configure the RADIUS server with a pre-shared key.

Configure all wireless workstations with client certificates.

Explanation:

When you use 802.1x authentication for wireless networks, a RADIUS server is implemented to centralize authentication. A centralized authentication database is used to allow wireless clients to roam between cells but authenticate using the same account information. To authenticate devices, PKI is required to issue certificates. At a minimum, the RADIUS server must have a server certificate. However, to support mutual authentication, each RADIUS client must also have a certificate. Remember that in a RADIUS solution, each wireless access point is a RADIUS client, not the wireless devices. The wireless access points forward the credentials from wireless devices to the RADIUS server for authentication.

Pre-shared keys are not used for authentication in an 802.1x solution.

q_wl_auth_access_radius_04_secp8

You are replacing a wired business network with an 802.11g wireless network. You currently use Active Directory on the company network as your directory service. The new wireless network has multiple wireless access points, and you want to use WPA2 on the network.

What should you do to configure the wireless network? (Select two.)

Answers:

***Install a RADIUS server and use 802.1x authentication**

***Configure devices to run in infrastructure mode**

Configure devices to run in ad hoc mode

Use open authentication with MAC address filtering

Use shared secret authentication

Explanation:

When using wireless access points, configure an infrastructure network. Because you have multiple access points and an existing directory service, you can centralize authentication by installing a RADIUS server and using 802.1x authentication.

Use ad hoc mode when you need to configure a wireless connection between two hosts.

Use open authentication with WEP or when you do not want to control access to the wireless network.

Use shared secret authentication with WPA or WPA2 when you can't use 802.1x.

q_wl_auth_access_security_secp8

You've just finished installing a wireless access point for a client.

What should you do to prevent unauthorized users from using the access point (AP) configuration utility?

Answers:

***Change the administrative password on the AP.**

Isolate the AP from the client's wired network.

Change the channel used by the AP's radio signal.

Implement MAC address filtering.

Explanation:

You should change the administrative password used by the AP. Many AP manufacturers use a default administrative username and password that are well known. If you don't change these parameters, anyone connecting to the AP can easily guess the password required to access the AP's configuration.

A wired and wireless network are already isolated from each other.

Changing the channel will not prevent unauthorized users from finding the channel and prevent them from accessing the available Wi-Fi.

MAC address filtering allows you to block traffic coming from certain known machines or devices. However, in this scenario, you are attempting to block unauthorized incoming traffic.

q_wl_auth_access_wpa2_secp8

You need to configure a wireless network using WPA2-Enterprise.

Which of the following components should be part of your design? (Select two.)

Answers:

TKIP encryption

***AES encryption**

WEP encryption

***802.1x**

Open authentication

Pre-shared keys

Explanation:

To configure WPA2-Enterprise, you need a RADIUS server to support 802.1x authentication. WPA2 uses AES for encryption.

WPA2-PSK, also called WPA2-Personal, uses pre-shared keys for authentication. WPA uses TKIP for encryption.

Using open authentication, any wireless device can authenticate with the access point, but the device can communicate only if its Wired Equivalent Privacy (WEP) keys match the access point's WEP keys. In this scenario, you are using WPA (not WEP) for authentication.

8.7 Data Transmission Security

As you study this section, answer the following questions:

How does SSL verify authentication credentials?

What protocol is the successor to SSL 3.0?

How can you tell that a session with a web server is using SSL?

What is the difference between HTTPS and S-HTTP?

What does it mean when HTTPS is stateful?

What is the difference between IPsec tunnel mode and transport mode?

In this section, you will learn to:

Add TLS to a website

Allow SSL connections

Require IPsec for communications

The key terms for this section include:

Term	Definition
Secure Sockets Layer (SSL)	A protocol that secures messages being transmitted on the internet.
Transport Layer Security (TLS)	A protocol that secures messages being transmitted on the internet. It is the successor to SSL 3.0.
Secure Shell (SSH)	A protocol that allows for secure interactive control of remote systems.
Hyper Text Transfer Protocol Secure (HTTPS)	A secure form of HTTP that uses either SSL or TLS to encrypt sensitive data before it is transmitted.
Secure Hypertext Transfer Protocol (S-HTTP)	An alternate protocol that is not widely used because it is not as secure as HTTPS.
Internet Protocol Security (IPsec)	A set of protocols that provides secure data transmission over unprotected TCP/IP networks.
Authentication Header (AH)	A protocol within IPsec that provides authenticity, non-repudiation, and integrity.
Encapsulating Security Payload (ESP)	A protocol within IPsec that provides all the security of AH plus confidentiality.
Security Association (SA)	The establishment of shared security information between two network entities to support secure communications.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions. Encryption Transport/communication Asymmetric Key exchange Certificates Certificate authorities 2.4 Given a scenario, analyze indicators of malicious activity. Network attacks Domain Name System (DNS) attacks 3.2 Given a scenario, apply security principles to secure enterprise infrastructure. Secure communication/access Virtual private network (VPN) Remote access Tunneling Transport Layer Security (TLS) Internet Protocol Security (IPSec)
TestOut Security Pro	3.2 Implement application defenses 3.2.3 Configure web application security 4.2 Implement Encryption Technologies 4.2.1 Encrypt data communications

8.7.1 Secure Protocols (Lesson Video)

Transcript:

FILE NOT FOUND

8.7.2 Secure Protocol Facts

Many protocols created in the past were designed with few to no security controls. An unsecured protocol is one that does not provide authentication or encryption or one that uses plaintext for passing authentication information or data. Newer protocols with security controls include Secure Sockets Layer (SSL), Transport Layer Security (TLS), Secure Shell (SSH), HyperText Transfer Protocol (HTTP), and HyperText Transfer Protocol Secure (HTTPS).

This lesson covers the following topics:

- Remote desktop
- Secure Shell
- HTTPS and S-HTTP

Remote Desktop

A remote access VPN joins the user's PC or smartphone to a remote private network via a secure tunnel over a public network. Remote access can also be a means of connecting to a specific computer over a network. This type of remote access involves connecting to a terminal server on a host using software that transfers shell data only. The connection could be a client and terminal server on the same local network or across remote networks.

A graphical remote access tool sends screen and audio data from the remote host to the client and transfers mouse and keyboard input from the client to the remote host. Microsoft's Remote Desktop Protocol (RDP) can be used to access a physical machine on a one-to-one basis.

Alternatively, a site can operate a remote desktop gateway that facilitates access to virtual desktops or individual apps running on the network servers. RDP connections are encrypted by default. There are several popular alternatives to Remote Desktops. Most support remote access to platforms other than Windows (macOS and iOS, Linux, Chrome OS, and Android, for instance). Examples include TeamViewer (teamviewer.com/en) and Virtual Network Computing (VNC) , which is implemented by several different providers (notably realvnc.com/en).

In the past, these remote desktop products required a dedicated client app. Remote desktop access can now just use a web browser client. The canvas element introduced in HTML5 allows a browser to draw and update a desktop with relatively little lag. It can also handle audio. This is referred to as an HTML5 VPN or as a clientless remote desktop gateway (guacamole.apache.org). This solution uses a protocol called WebSocket, which enables bidirectional messages to be sent between the server and client without requiring the overhead of separate HTTP requests.

Secure Shell

Secure Shell (SSH) is the principal means of obtaining secure remote access to a command line terminal. The main uses of SSH are for remote administration and secure file transfer (SFTP). Numerous commercial and open - source SSH products are available for all the major NOS platforms. The most widely used is OpenSSH (openssh.com).

SSH servers are identified by a public/private key pair that is referred to as the Host Key. Hostnames can be mapped to Host Keys manually by each SSH client or through various enterprise software products designed for SSH Host Key management.

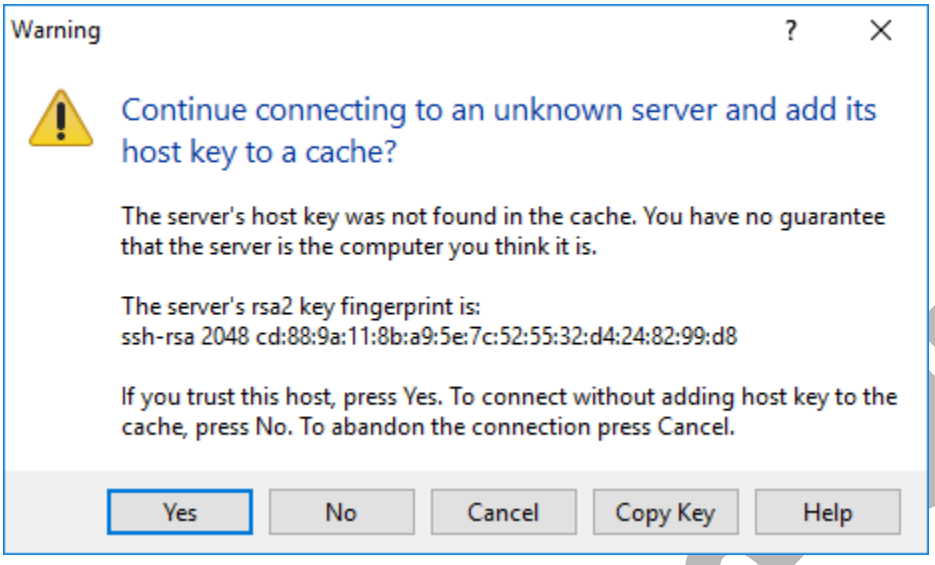

Confirming the SSH server's Host Key using the PuTTY SSH client (Screenshot used with permission from PuTTY.)

The Host Key must be changed if any compromise of the host is suspected. If an attacker has obtained the private key of a server or appliance, they can masquerade as that server or appliance and perform a spoofing attack, usually with a view to obtaining other network credentials.

The server's Host Key is used to set up a secure channel for the client to submit authentication credentials.

Type	Description
SSH client authentication	SSH allows various methods for the client to authenticate to the server. Each of these methods can be enabled or disabled as required on the server using the /etc/ssh/sshd_config file: **Username/password** - is when the client submits credentials verified by the SSH server either against a local user database or using a RADIUS server. **Public key authentication** - is when each remote user's public key is added to a list of keys authorized for each local account on the SSH server. **Kerberos** - is when the client submits the Kerberos credentials (a Ticket Granting Ticket) obtained when the user logs onto the workstation to the server using the Generic Security Services Application Program Interface (GSSAPI). The SSH server contacts the Ticket Granting Service (in a Windows environment, this will be a domain controller) to validate the credentials. Managing valid client public keys is a critical security task. Many recent attacks on web servers have exploited poor key management. If a user's private key is compromised, delete the public key from the appliance, regenerate the key pair on the user's (remediated) client device, and copy the public key to the SSH server. Always delete public keys when the user's access permissions have been revoked.

SSH commands	SSH commands are used to connect to hosts and set up authentication methods. To connect to an SSH server at 10.1.0.10 using an account named "bobby" and password authentication, run: ssh bobby@10.1.0.10 The following commands create a new key pair and copy it to an account on the remote server: ssh-keygen -t rsa ssh-copy-id bobby@10.1.0.10 At an SSH prompt, you can now use the standard Linux shell commands. Use exit to close the connection. You can use the scp command to copy a file from the remote server to the local host: scp bobby@10.1.0.10:/logs/audit.log audit.log Reverse the arguments to copy a file from the local host to the remote server. To copy the contents of a directory and any subdirectories (recursively), use the -r option.

HTTPS and S-HTTP

A common unsecured protocol is the HyperText Transfer Protocol (HTTP). HTTP is used for exchanging web content and passes data in cleartext. HTTP uses TCP port 80 and is stateless, which means, by default, it does not keep track of clients. To solve this problem, cookies can be used to keep track of the client's behavior. To secure HTTP, use one of the following protocols:

Protocol	Description
HTTPS	HyperText Transfer Protocol Secure is a secure form of HTTP that uses SSL or TLS to encrypt sensitive data before it is transmitted. HTTPS: Is stateful, which means that it keeps track of the client. To do this, the client must communicate with the same HTTPS server for the duration of the session. Load balancing is not possible during the connection and is available only to initially determine which server will handle the client's session. Requires TCP port 443 inbound on the webserver to be open. Can be identified by verifying that the URL starts with https:// or by looking for a lock symbol in the browser. Double-clicking on the lock icon displays the certificate.
S-HTTP	Secure HyperText Transfer Protocol (S-HTTP) is an alternate protocol that is not widely used because it is not as secure as HTTPS. S-HTTP: Is connectionless, unlike SSL, which is connection-oriented.

	Provides only message security, unlike HTTPS, which provides a full secure channel for all messages. Does not use port 443.

8.7.3 Add TLS to a Website (Demo Video)

Transcript:

In this demonstration, we'll look at using SSL for a website. We're going to use Windows Server 2022 to configure a website. We're on this client machine in Windows 11 with a browser open.

As you can see, we're using HTTP for this website, so traffic sent between the client and the website is unsecured. To secure communication between the client and the website, we can add SSL to the website. We do this by editing the bindings. The bindings identify the protocols that are used on the website. Currently, we support the HTTP protocol, which uses port 80. To use SSL, we need to add bindings for HTTPS.

Before we get into this demo, we need to talk about the terminology used for SSL and TLS. For the last few years, there's been a large push in the IT industry to deprecate TLS versions 1.0 and 1.1 and require all encrypted internet traffic to use TLS version 1.2 or later. SSL is the predecessor to TLS. The last version of SSL, version 3.0, was found to be insecure. SSL 3.0 was deprecated in 2015 by the IETF in RFC 7568.

If SSL was deprecated in 2015, why do we still see the SSL acronym everywhere? You still see websites selling SSL certificates, and many of Microsoft's own tools still reference SSL certificates—including tools you'll see in this demo. The answer is that because the SSL term was so prevalent at the time it was deprecated, people are still adjusting to using TLS. When you hear someone talk about or see documentation referencing SSL, they are usually talking about TLS. The SSL protocol hasn't been used on the internet for years and isn't supported in modern web browsers. Up-to-date versions of Chrome, Edge, and Firefox require TLS connections for encryption and won't allow you to connect to a web server that's trying to use SSL.

Please make a note of this as we go through this demo and as you get into your IT career. While the term SSL is still commonly used in both speech and documentation, the protocol SSL has been fully replaced by TLS for many years.

We're going to go to the Internet Information Service Manager, so let's switch over to our server. Let's go to Start and Server Manager. In the tools, we'll select IIS Manager. First, we're going to go down to our default website. We can right-click the Default Web Site we've created. Go to edit Bindings. We can see our HTTP port 80 here. We'll click Add.

We can switch this to HTTPS. You can see that it automatically switches to port 443. Notice that secure communications with this website use port 443 instead of port 80. When using SSL, you must also have a server certificate, or SSL certificate. The certificate identifies the server and validates its identity. In this case, we don't currently have a certificate set up, so we're going to click Cancel and Close. We'll go to our server and go to Server Certificates.

Now, if our website were available on the internet, we'd create a certificate request here and send that request to a well-known public key infrastructure, or PKI, such as DigiCert. When that request was returned, we'd use this complete certificate request option to import the completed certificate onto our server. In this demonstration, we're going to use Create Self-Signed Certificate. The self-signed certificate gives us a certificate that we can use. However, clients won't automatically trust the certificate that we use. We'll see how this works later. We'll enter a name for the certificate. And we'll call it Cert 2022 and click OK.

You can see our certificate here. With this certificate created, we'll go back to our Default Web Site, right-click, Edit Bindings > Add > HTTPS. It switches to the 443. HTTP/2 is how semantics flow over TCP connections. Most browsers support HTTP/2, so disabling this may not be needed. OCSP Stapling checks to see if a certificate is revoked and has been an alternative to the Certificate Revocation List (CRL) Protocol. We'll leave both of those unchecked for now. We're going to go down, and we can see our Cert 2022 right here. Select it and click OK and Close.

That certificate's created. Now, we can go over to our client machine and test this out. We're back on our Windows 11 client machine. We need to put https:// in front of the address and hit enter. The first message we get is that there's a problem with the website's security certificate. We started to use HTTPS, but our computer recognized that the certificate wasn't issued by a trusted certificate authority. In other words, the server has no proof, other than its own words, to verify its identity.

In this case, we know who the server is, so we can continue to the website and accept the certificate. Notice that we're accessing the website. The URL is HTTPS. Let's go ahead and click Continue. The message told us there was a problem with the certificate, but we're still using HTTPS to connect to the website. In this case, we're giving the user the option of using HTTPS or HTTP when communicating with our server. A more likely example is to require SSL for certain communications.
In this case, we're going to require SSL for the entire website, so a connection won't be granted unless the client and the server are both using SSL. To do this, we're going to go back to our web server, to Internet Information Services Manager.
We're going to click the Default Web Site. Go to SSL Settings. We're going to tell it to Require SSL. We'll check the box and hit Apply to save those changes. Now, we can go back to the client.
First, let's try to access the website using just HTTP. We'll just go ahead and take out the S. That's all we have to do. Hit Enter. We get this error that says access is denied. Now, let's go ahead and put the S back. It's basically denying us unless we're using an SSL certificate with a secure connection. You can see it went ahead and let us back in. The website requires SSL.
We were trying to use HTTP, but it won't let us. Now that we're using HTTPS, we can access the website because the client and the server are both using SSL. In this example, we've required SSL on the entire website, meaning that any communication to the website must use SSL or be denied. It's unlikely that you may require SSL only for specific parts of the website.
Today, it's recommended always to use HTTPS with a public-facing website, regardless of whether it's a basic website or one used for credit card purchases. By only using HTTP, hackers have easy access to view traffic in plain text.
That's it for this demo. We've gone over how to configure SSL on a website using Windows Server 2022. We've also clarified that SSL as a protocol is no longer used and has been fully replaced by TLS, even though the term SSL is still widely used in common speech and documentation.

8.7.4 Allow SSL Connections (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You currently run a website on the CorpWeb server. You want to allow SSL connections to this website.

In this lab, your task is to add a binding to the CorpNet website using the following settings:

Website: **www.corpnet.xyz**

Protocol: **HTTPS**

Port: **443**

SSL certificate: **www.corpnet.xyz**

Explanation

Complete this lab as follows:

Open the IIS Manager to the CorpNet.xyz site.

From the Server Manager's menu bar, select **Tools** > **Internet Information Services (IIS) Manager** .

Expand **CorpWeb (CorpNet.com\Administrator)** > **Sites** .

Select **CorpNet.xyz** .

Add a binding to the CorpNet website.

From the Actions pane (far right), select **Bindings** .

Select **Add** .

Using the *Type* drop-down menu, select **HTTPS** .

Make sure the port is set to **443** .

Using the *SSL certificate* drop-down menu, select **www.CorpNet.xyz** and then select **OK** .

Select **Close** .

8.7.5 IPsec (Lesson Video)

Transcript:

In this lesson, we'll cover Internet Protocol Security, or Ipsec, which is a set of security protocols used to provide authentication and asymmetric encryption for TCP/IP network traffic. It's also used in VPNs.

IPsec operates on top of the Network layer, or Layer 3, of the OSI model. The OSI model provides standards for communication in a computing system, which allows interoperability of various systems regardless of their internal structures. It's composed of seven layers, each layer being served by the layer below it. The bottom three layers are the media layers, and the top four layers are the host layers. Layer 3's primary goal is communication, not security. Applying IPsec on top of Layer 3 provides security to the communication through mutual authentication, integrity, non-repudiation, and confidentiality.
The first protocol in the IPsec suite is called Authentication Header, or AH. It encapsulates all the Host layer information, which includes Layers 4 through 7, and replaces it with an authentication header. AH will detect if the IP address on Layer 3 changes through a checksum for all the data from Layers 3 through 7. This provides authenticity and integrity, but not confidentiality because the encapsulated data in the packet is not encrypted. If the packet were intercepted by a sniffer, the data could be read.
AH provides authentication information in the form of a keyed hash, which is based on all the bytes in the packet. AH authenticates packets by digitally signing them, which prevents replay attacks. AH's assigned IP protocol number is 51.
The second IPsec protocol option is Encapsulating Security Protocol, or ESP. It's commonly used with IPsec because it provides everything that AH provides plus confidentiality. ESP's assigned IP protocol number is 50.
Like AH, ESP encapsulates the Host layers, 4 through 7, into a new Layer 4 header called an ESP header. Unlike AH, it encrypts the encapsulated data, preventing someone from reading the data if it's sniffed. Also unlike AH, ESP doesn't notice when a lower level IP address changes because there isn't a checksum that includes Layer 3 information. To solve this problem, ESP works best with Network Address Translation, or NAT.
An important consideration is that NAT can have problems with IPsec. Problems can come because IPsec secures the headers of packets and detects if the packets have been tampered with, and NAT needs to tamper with packets by changing source and destination IP addresses and ports.
To fix this, something called NAT Traversal, or NAT-T, was created. NAT-T is designed to allow IPsec to function properly through a NAT device. It does this by encapsulating ESP packets inside a UDP packet and uses UDP port 4500.
Another part of IPsec to consider is the set of specifications that negotiate between nodes to establish the IPsec relationship called a Security Association, or SA. These specifications can include cryptographic keys, authentication preferences, certificates, and algorithm selections.
For example, both endpoints can agree to use the SHA-1 hashing algorithm instead of MD5, AES for symmetric encryption, and RSA for asymmetric encryption. Each tunnel of data uses three different security associations.
First, a management channel is established so that routers or network nodes can exchange security information. Then an outbound security association and an inbound security association are established, each with a unique identifier that's

included with each packet sent across the channel. A security association can be established manually or automatically through a protocol called IKE, Internet Key Exchange.
IKE helps establish automatic SAs and a secure tunnel by providing a protected exchange of keys before the full IPsec transmission begins. IKE uses a Diffie-Hellman key exchange to establish a shared session secret. Mutual authentication is provided either by pre-shared keys on both endpoints or through certificates issued by a CA. IKE can also help automate the selection of the best security association for each connection.
For example, if both endpoints support Triple DES and AES for symmetric encryption, then IKE can help them negotiate the strongest security method, AES. Once that's agreed on, they're placed into their SAs. Now their symmetric encryption algorithm is set to encrypt and decrypt data with the IPsec tunnel.
Finally, there are two modes that are available to use with IPsec circuits, including VPNs.
The first is called tunnel mode. It's commonly used to provide secure communication between two network gateways. For example, it will allow any user at Corporate Office A to communicate safely with any user at Corporate Office B because both offices have secure gateways acting as IPsec proxies.
The second is called transport mode. It's used when two hosts connect directly with each other and the circuit is broken off after the session ends.
That's it for this lesson. In this lesson, we discussed how IPsec secures IP traffic across the network. Then we looked at elements of the IPsec suite. First, we looked at AHs, or Authentication Headers. Then we looked at ESP, or Encapsulating Security Payload. Remember to use NAT-T to enable ESP to function properly. Next, we discussed how Security Association specifications work. Then we looked at using Internet Key Exchange protocol to implement the SAs. Finally, we covered the two IPsec modes, tunnel mode and transport mode. Tunnel mode allows protected communication between users at both ends of the secure tunnel through their gateways, and transport mode is for direct connection between hosts for a specific session.

8.7.6 IPsec Facts

This lesson covers the following topics:

- Remote access architecture
- Internet Protocol Security tunneling
- Transport layer security tunneling

Remote Access Architecture

Remote access networking means that the user's device does not make a direct cable or wireless connection to the network. The connection occurs over or through an intermediate network.

Historically, remote access used analog modems connecting over the telephone system. These days, most remote access is implemented as a virtual private network (VPN) , running over Internet Service Provider (ISP) networks.

With a remote access VPN, clients connect to a VPN gateway on the edge of the private network. This client-to-site VPN topology is the "telecommuter" model, allowing homeworkers and employees working in the field to connect to the corporate network. The VPN protocol establishes a secure tunnel to keep the contents private, even when the packets pass over ISPs' routers.

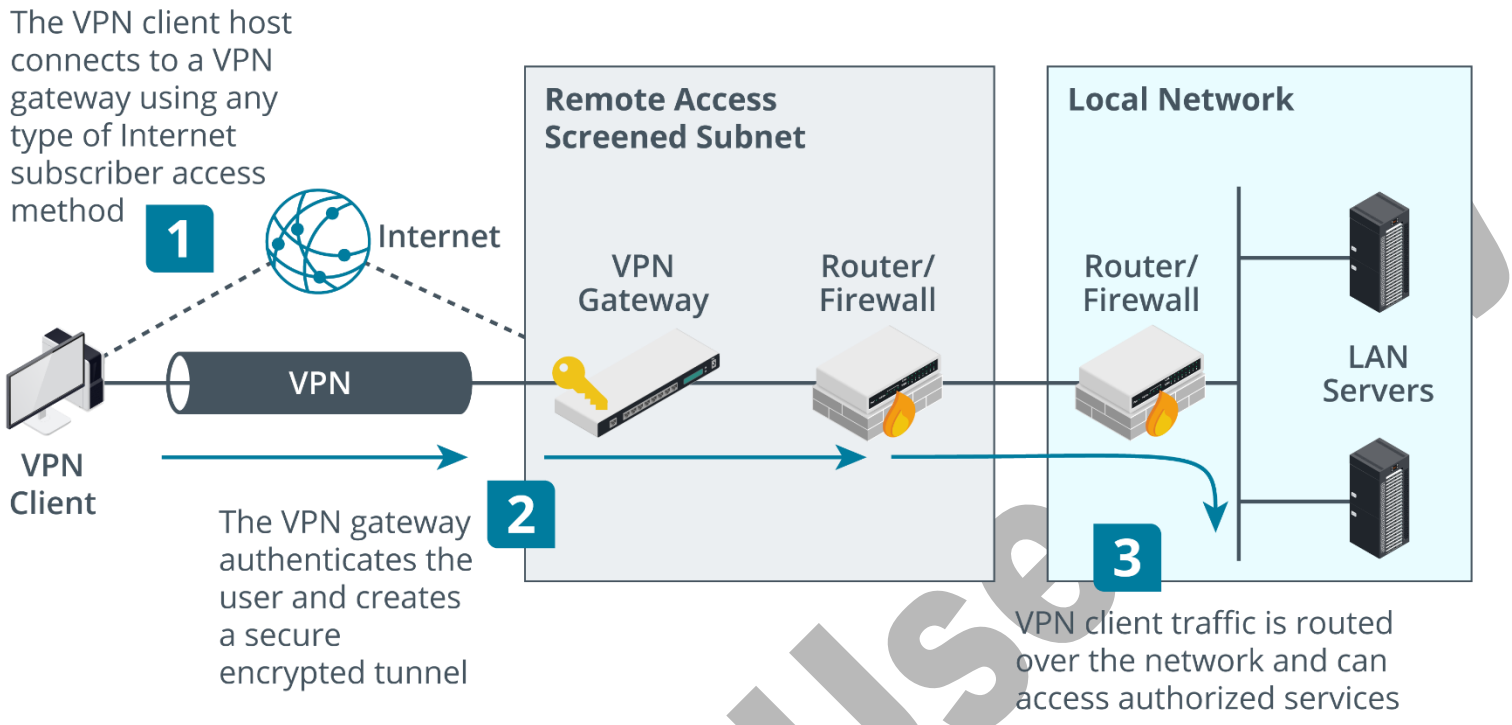

Remote access VPN. (Images © 123RF.com.)

A VPN can also be deployed in a site-to-site model to connect two or more private networks. Whereas remote access VPN connections are typically initiated by the client, a site-to-site VPN is configured to operate automatically. The gateways exchange security information using whichever protocol the VPN is based on. This establishes a trust relationship between the gateways and sets up a secure connection through which to tunnel data. Hosts at each site do not need to be configured with any information about the VPN. The routing infrastructure at each site determines whether to deliver traffic locally or send it over the VPN tunnel.

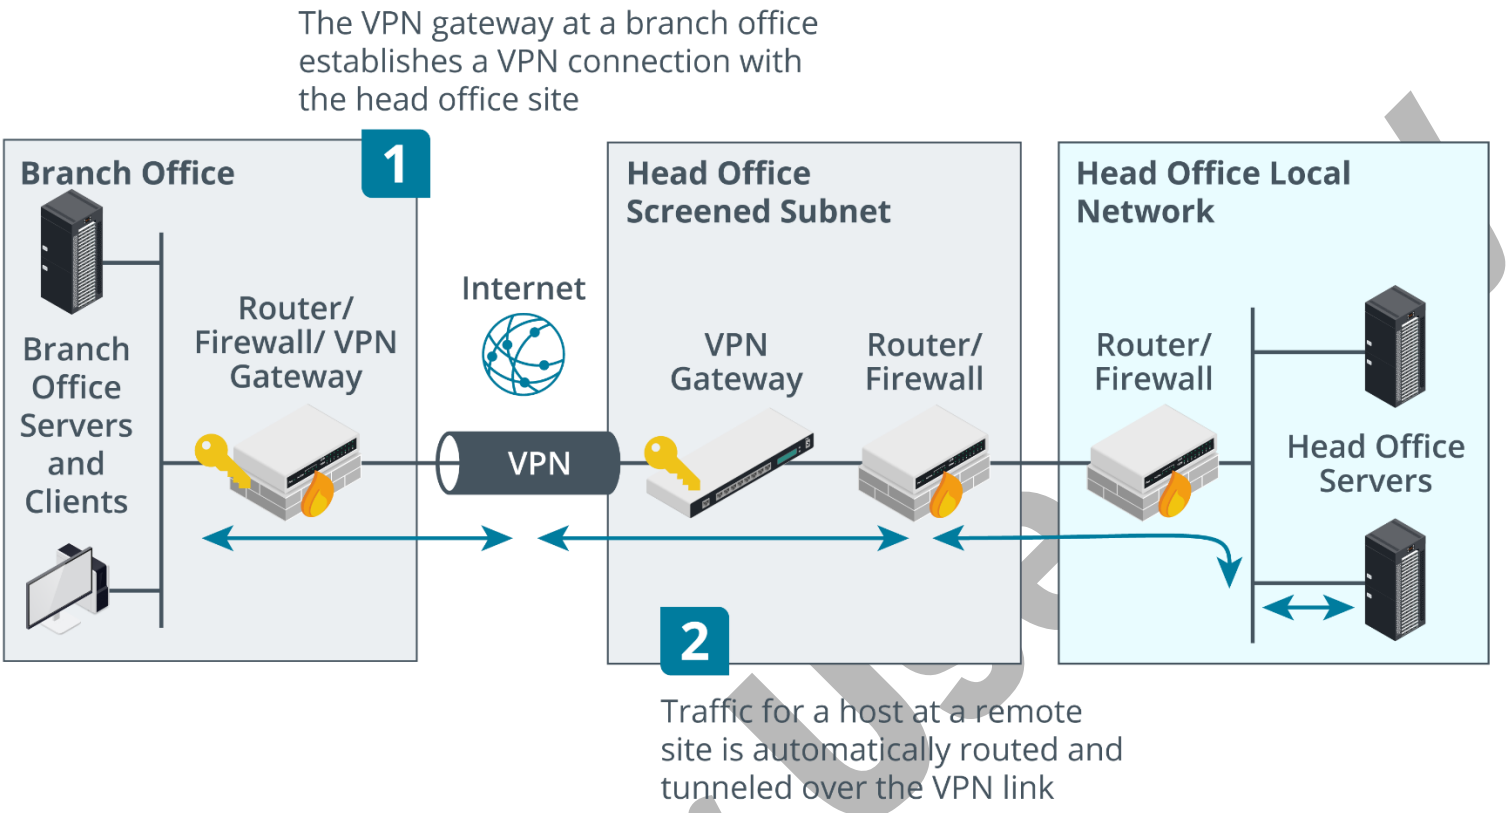

Site-to-site VPN. (Images © 123RF.com.)

A third topology is a **host-to-host tunnel** . This is a means of securing traffic between two computers where the private network is not trusted.

Several VPN protocols have been used over the years. Legacy protocols, such as the Point-to-Point Tunneling Protocol (PPTP) , have been deprecated because they do not offer adequate security. Transport Layer Security (TLS) and Internet Protocol Security (IPsec) are now the preferred options for configuring VPN access.

Transport Layer Security Tunneling

A transport layer security (TLS) VPN means the client connects to the remote access server using digital certificates. The server certificate identifies the VPN gateway to the client. Optionally, the client can also be configured with its own certificate. This allows for mutual authentication, where both server and client prove their identity to one another. TLS creates an encrypted tunnel for the user to submit authentication credentials. These would normally be processed by a RADIUS server. Once the user is authenticated and the connection is fully established, the VPN gateway tunnels all communications for the local network over the secure socket.

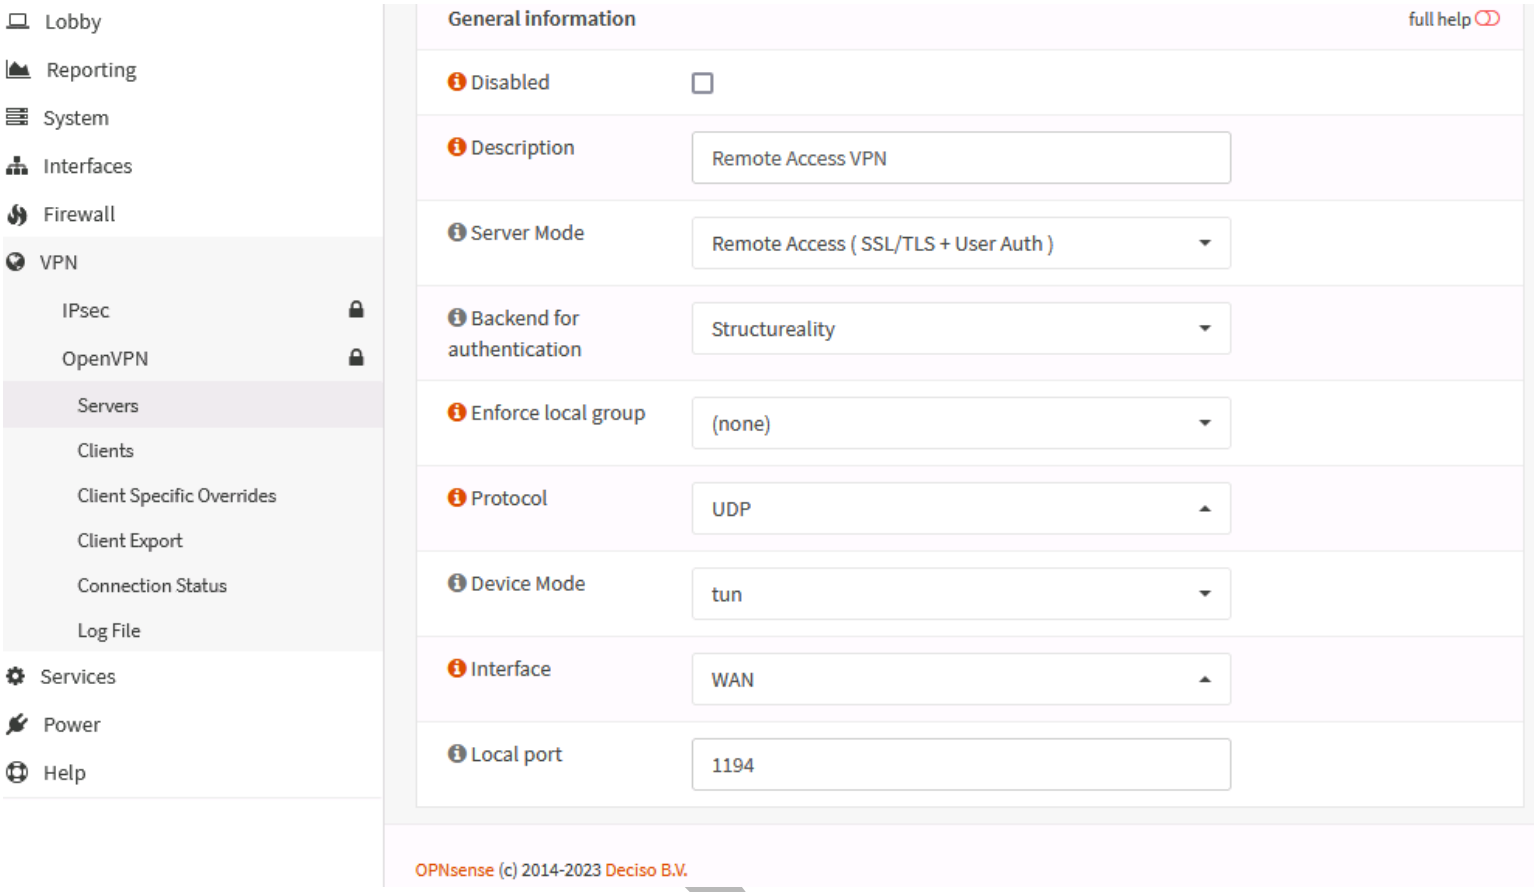

Configuring an OpenVPN server in the OPNsense security appliance. This configuration creates a remote access VPN. Users are authenticated via a RADIUS server on the local network. (Screenshot courtesy of OPNsense.)

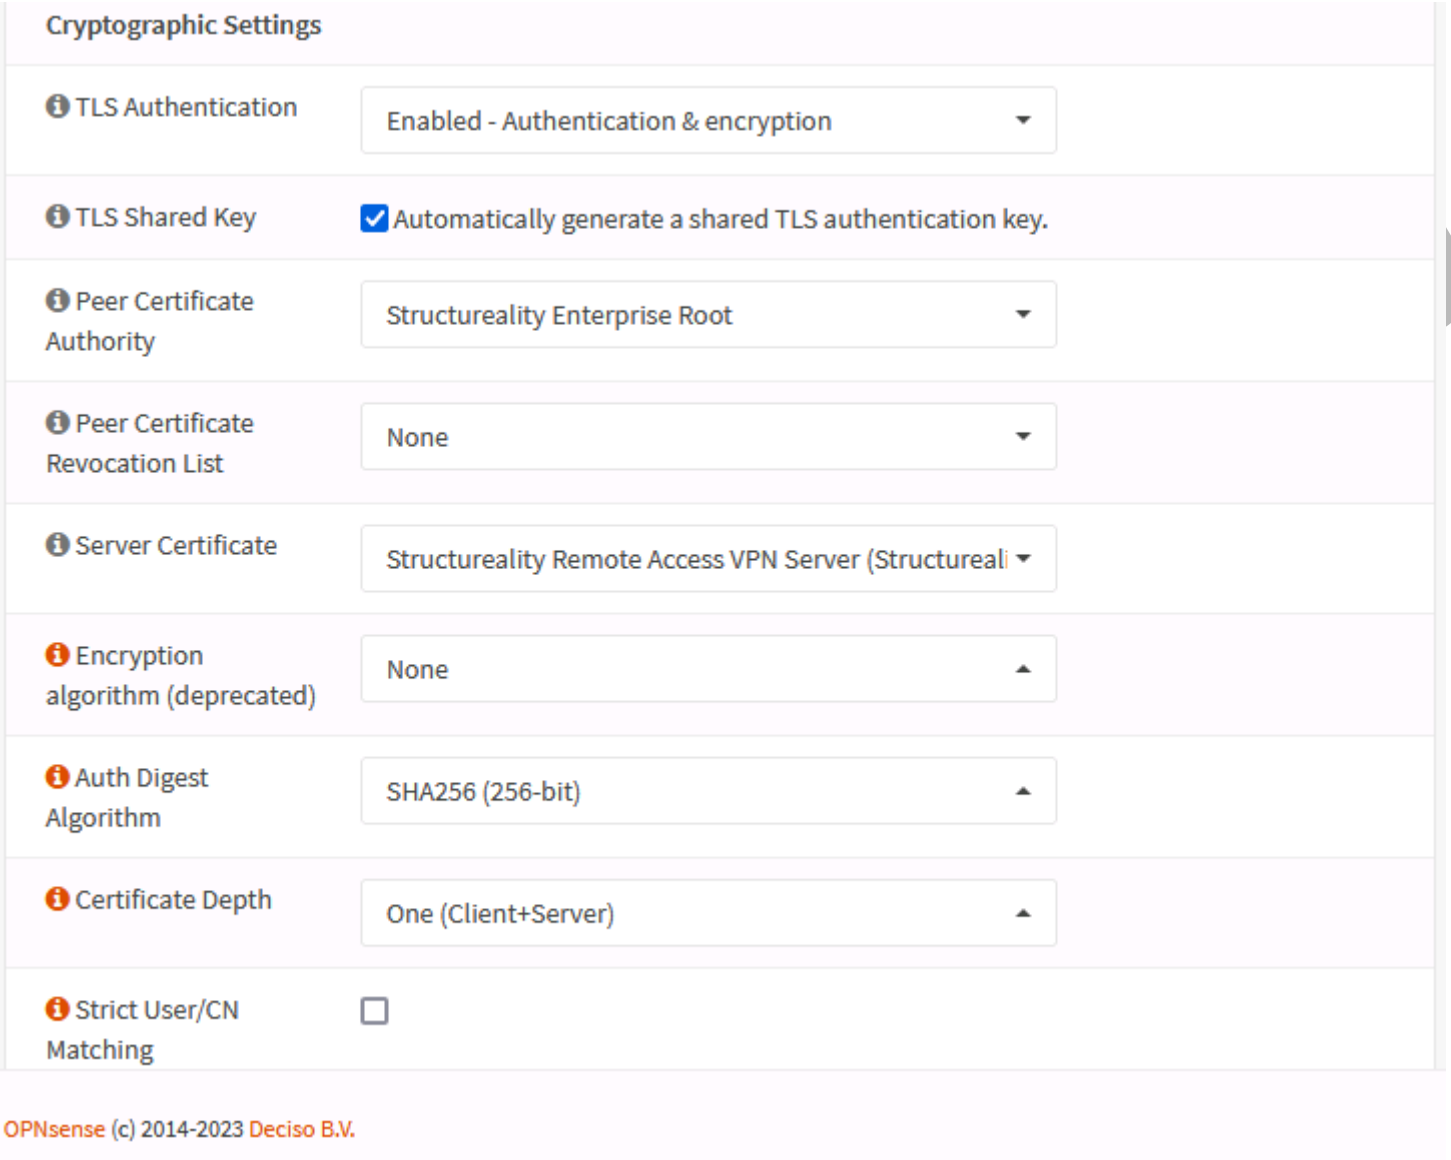

Configuring a server certificate for OpenVPN in the OPNsense security appliance. (Screenshot courtesy of OPNsense.)

A TLS VPN can use either TCP or UDP. UDP might be chosen for marginally superior performance, especially when tunneling latency-sensitive traffic such as voice or video. TCP might be easier to use with a default firewall policy. TLS over UDP is also referred to as Datagram TLS (DTLS).

It is important to use a secure version of TLS. The latest version at the time of writing is TLS 1.3. TLS 1.2 is also still supported. Versions earlier than this are deprecated.

Internet Protocol Security Tunneling

Transport Layer Security is applied at the application level. Internet Protocol Security (IPsec) operates at the network layer of the OSI model (layer 3). This means that it can be implemented without having to configure specific application support and that it incurs less packet overhead.

There are two core protocols in IPsec, which can be applied singly or together, depending on the policy:

Authentication Header (AH) - performs a cryptographic hash on the whole packet, including the IP header, plus a shared secret key (known only to the communicating hosts), and adds this value in its header as an Integrity Check Value (ICV). The recipient performs the same function on the packet and key and should derive the same value to confirm that the packet has not been modified. The payload is not encrypted, so this protocol does not provide confidentiality.

Encapsulating Security Payload (ESP) can be used to encrypt the packet rather than simply calculating an ICV. ESP attaches three fields to the packet: a header, a trailer (providing padding for the cryptographic function), and an Integrity Check Value. Unlike AH, ESP excludes the IP header when calculating the ICV.

IPsec can be used in two modes:

Transport mode - is used to secure communications between hosts on a private network. When ESP is applied in transport mode, the IP header for each packet is not encrypted, just the payload data. If AH is used in transport mode, it can provide integrity for the IP header.

IPsec datagram using AH and ESP in transport mode.

Tunnel mode - is used for communications between VPN sites across an unsecured network. With ESP, the whole IP packet (header and payload) is encrypted and encapsulated as a datagram with a new IP header. AH has no use case in tunnel mode, as confidentiality is usually required.

IPsec datagram using ESP in tunnel mode.

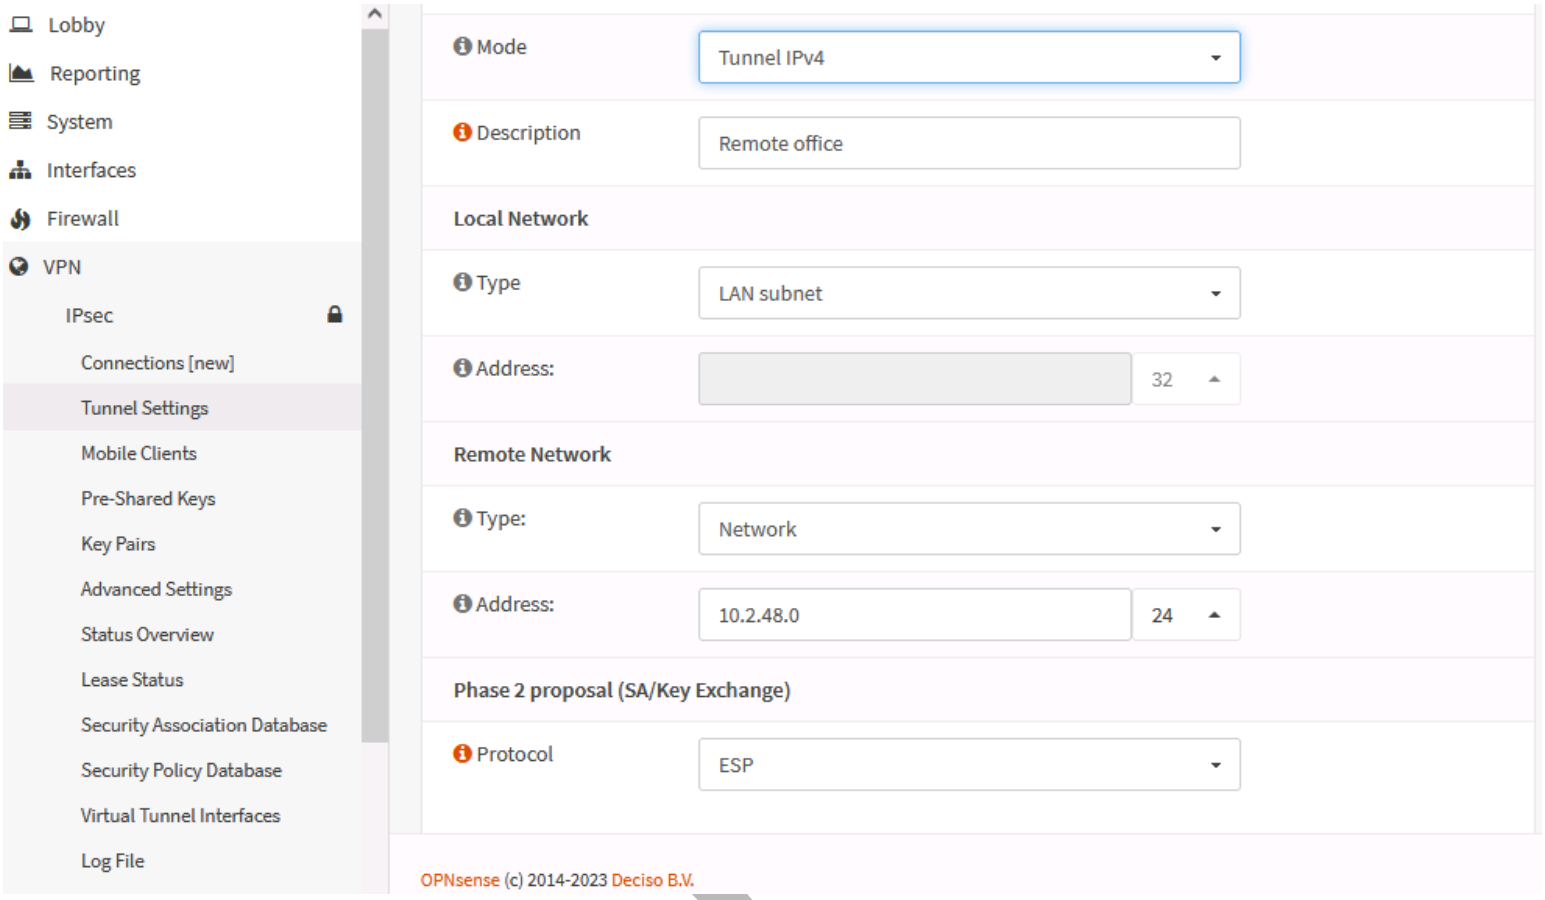

Configuring a site-to-site VPN using IPsec tunneling with ESP encryption in the OPNsense security appliance. (Screenshot courtesy of OPNsense.)

Each host or router that uses IPsec must be assigned a policy. An IPsec policy sets the authentication mechanism and the use of AH/ESP and transport or tunnel mode for a connection between two peers.

IPsec's encryption and hashing functions depend on a shared secret. The secret must be communicated to both peers, who must perform mutual authentication to confirm one another's identity. The Internet Key Exchange (IKE) protocol implements an authentication method, selects which cryptographic ciphers are mutually supported by both peers and performs key exchange. The set of properties is referred to as a security association (SA).

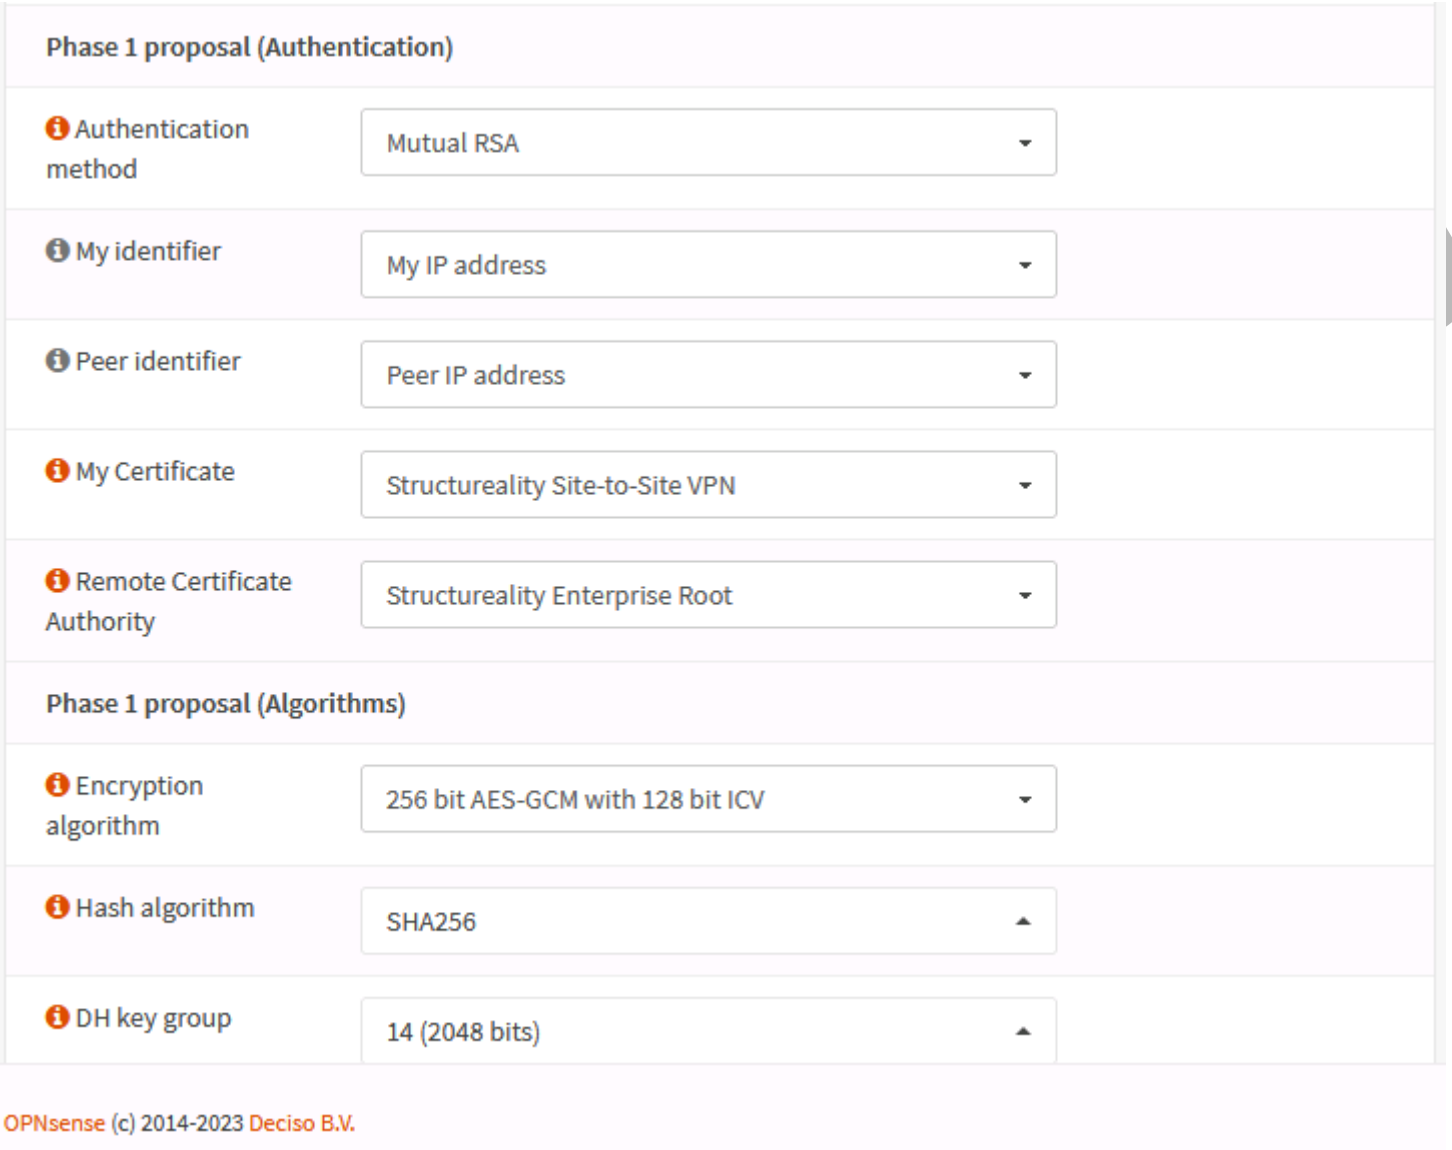

Configuring IKE for certificate-based authentication in the OPNsense security appliance. (Screenshot courtesy of OPNsense.)

IKE negotiations take place over two phases:

Phase I establishes the identity of the two peers and performs key agreement using the Diffie-Hellman algorithm to create a secure channel. Two methods of authenticating peers are commonly used:

Digital certificates - are issued to each peer by a mutually trusted certificate authority to identify one another.

Pre-shared key (group authentication) - is when the same passphrase is configured on both peers.

Phase II uses the secure channel created in Phase I to establish which ciphers and key sizes will be used with AH and ESP in the IPsec session.

There are two versions of IKE. Version 1 was designed for site-to-site and host-to-host topologies and required a supporting protocol to implement remote access VPNs. IKEv2 has some additional features that have made the protocol popular for use as a stand-alone remote access client-to-site VPN solution. The main changes are the following:

Supports EAP authentication methods, allowing, for example, user authentication against a RADIUS server.

Provides a simple setup mode that reduces bandwidth without compromising security.

Allows network address translation (NAT) traversal and MOBIKE multihoming. NAT traversal makes it easier to configure a tunnel allowed by a home router/firewall. Multihoming means that a smartphone client with Wi-Fi and cellular interfaces can keep the IPsec connection alive when switching between them.

8.7.7 Requiring IPsec for Communications (Demo Video)

Transcript:

In this demonstration, we'll look at requiring IPsec for communications between two network hosts. We'll use a Windows server and a Windows workstation.

To secure communications with IPsec, we use the Windows Defender Firewall with Advanced Security. Let's search for Windows Defender Firewall. There it is. Let's expand this. You'll notice that no rules are defined in the connection security rules here, so we need to define a new rule.
To define a new rule, you right-click Connection Security Rules and create a New Rule. The wizard displays. First, we choose the type of rule to use for IPsec. We have several options. We could configure an Isolation rule, which restricts connections based on authentication criteria such as domain membership or health status. We have an Authentication exemption, which will prevent authenticated connections with specific computers, in other words, it exempts computers from connecting.
We have Server-to-server rules. These authenticate connections between specified computers. Tunnel rules authenticate connections between two computers; it creates a tunnel between the two. We also can create Custom rules.
We want to set up an Isolation rule. We'll select Isolation. Click Next. Now, we need to specify how to enforce authentication. You can either request to use IPsec or require the use of IPsec. If you request, the default option is the computer that attempts to use IPsec. If either computer in the communication channel can't use IPsec, it fails over to unsecured communications.
Notice that this rule is applied to both inbound and outbound connections. If this rule isn't secure enough, then you can increase security by requiring IPSec for inbound and outbound connections. We'll request authentication for inbound and outbound connections. Click Next, which is the default option. Now, we determine the type of authentication method we want to use.
The first option, the Default, is to use whatever authentication methods are specified in the default IPsec settings. Most of the time, you'll use computer-based Kerberos. To use Kerberos, both computers must be in the same Active Directory forest.
You also have the option to specify Computer and user Kerberos. Choosing this option also requires both computers to be in the same Active Directory forest. The systems must be re-authenticated using user authentication.
In this demo, we'll use Computer authentication. There's also an Advanced option. After you select Advanced, click Customize. Customize allows you to set up multiple authentication methods. If the first method we set up fails, the system defaults to the next available authentication type. On the left, we specify the first authentication method. If we click Add, the options are a computer using Kerberos, a computer using NTLM, or a certificate from a trusted certificate authority.
That last option is useful if both computers aren't in the same Active Directory forest. In this case, we can use certificates to authenticate both ends of the communication channel. We'll configure both computers to trust the issuing CA of the certificates used on the computers. We'll browse through to select the certificates that we want to use.
We also have the option of using a pre-shared key that would be used by both computers. It's noted here that this isn't recommended. A pre-shared key is the least secure option of these four. Let's go ahead and click Cancel.
On the left, we specify the first authentication method. If the authentication doesn't work, the system will use the second authentication method we configured here on the right. For this demo, both hosts are members of the same domain, so we'll use computer-based Kerberos authentication. Now we'll click Next.
Now, we can specify the network location profile to which this rule will be applied. Because we have a desktop workstation and a server that are members of the same domain and will always be attached to the same network, we can select domain. However, if you have a roaming computer such as a laptop, notebook, or something of that sort, this becomes a little more of an issue.

While at work, the notebook or laptop would use the domain profile. When the computer is used from home or another location, it would use a public or private profile on the system, depending on where we connect. For our purposes today, we'll use all three of these profiles. We could turn one or more of them off if we wanted to restrict that. Now, we need to name this rule. We'll enter Request IPsec policy. Click Finish. At this point, the rule has been configured, but we've configured only one-half of the communication channel.

We'll go to the workstation side and configure the same rule here. On the workstation side, we'll go to the Search box and type Windows Defender Firewall with Advanced Security. We'll go to Connection Security Rules and right-click. We'll configure this IPsec rule like we did on the server system. Click Isolation and then Next. Now, we'll click Request authentication for inbound and outbound connections. We'll go to Computer Kerberos and leave all three. We'll name it Request IPsec and click Finish.

From the workstation, let's initiate communications between these two computers. We'll go back to the Search box and open a command prompt. We'll ping the IP address of the server system. If we don't know the IP address of the server system, we can go over to the server system, open the Command Prompt, and run ipconfig.

You can see the server address here: 172.16.50.110. Back to our workstation, ping 172.16.50.110. We want this to be a continuous ping, so we'll cancel and restart it with the -t flag to continuously send pings.

You can see we're getting replies from that address. That's perfect. We've established that we can ping the server. The pinging process will continue, and we can monitor the connection between the two computers.

We can see the replies from the server, and it just keeps running. Now, let's switch back to the server and go back to Windows Defender Firewall with Advanced Security. We can see the connection rule. Let's go to the Monitoring tab so we can monitor the connection.

If we go to the Connection Security Rule, we can see that the Request IPsec rule is currently in progress. We'll go to Security Associations. Here, we can look inside the Main Mode or Quick Mode and see the connections between our Windows 11 machine and server.

We see information that the connection is working. There are two modes, the main mode and quick mode because IPsec tunnels are built in two different phases. Phase one is called main mode. Phase two is quick mode. Breaking the monitoring information into these two modes is useful. It allows you to troubleshoot each phase of the connection establishment process.

Here, in Main Mode, we see information about the connection. We have the two endpoints here. We can see that this address here is the Windows Server. This is the workstation address here and the remote address.

We can also see that we're using Kerberos authentication here. We're not requiring a second authentication at the moment. If we scroll over or look at this side, we can see the encryption in the data integrity mode.

Let's right-click and go to Properties. We can see a little more information about the connection. Let's click Cancel. We can do the same thing for Quick Mode and see the same information. Again, we can see both ends of the communication: a local address and a remote address. We can see that AH and ESP are components of quick mode.

Before we leave Windows Defender Firewall with Advanced Security, we want to show you the global settings.

Recall earlier when we set up our rule. We right-clicked Connection Security Rules. From there, we clicked New Rule, Request authentication, and then authentication method. The default is to use the authentication method specified in IPsec settings.

Now, we can configure IPsec settings in the Windows Firewall with Advanced Security. That way, we don't have to manually configure these settings each time we set up a rule. We can use the default option. Let's select Cancel. Now, we'll go to Windows Defender Firewall with Advanced Security, right-click, go to Properties, and then go to the IPsec tab.

Here, we configure the IPsec setting defaults. These are the settings applied unless you override them manually when you create the rules. Let's click Customize. We can configure how the key exchange works. Select Advanced, then select Customize, and specify the Security methods used. We can enter a specific algorithm to use or specify key lifetimes. These are settings for the main mode; data protection is for quick mode.

If we click Advanced on data protection from quick mode and click Customize on this screen, we can configure the AH and ESP algorithms for data integrity and encryption. Note that we have the algorithm listed first or the algorithm that's used first. Remember that integrity ensures that the data hasn't been tampered with in transit; that's hashing. Encryption encrypts the data so that it can't be sniffed and read. Let's click Cancel. We can configure the default authentication method. We can set that to Computer Kerberos if we want. We used that when we set up the rule. Let's click Cancel because we don't want to set these defaults right now. Click Cancel again.

The last thing we need to review is the tools you can use to manage IPsec communications from the command line. Let's go to the command prompt. We'll type netsh advfirewall and ?. It displays the commands we have available. We'll use the consec command. Let's press Up and run the consec command.

Here, you see the commands available within this context. We can add a new connection, delete all matching connection security rules, display the configuration script, display a list of commands, set new values for properties of an existing rule, and display a specified connection security rule.

That's it for this demonstration. In this demonstration, we used IPsec to secure communications between two network hosts. We used Windows Firewall with Advanced Security to set up connection security rules on both sides of the communication.
We also set up secured communications between a Windows server system and a Windows client system. We then used the ping command to establish communications between the two hosts. We monitored the connection to verify that IPsec was being used to secure communications between them. We then talked about configuring the IPsec defaults. We ended this demonstration by talking about how you can manage IPsec rules from the command line.

8.7.8 Modify Enterprise Capabilities to Enhance Security

8.7.9 Practice Questions (Section Quiz)

q_sec_prot_ftps_01_secp8

Which of the following is a secure alternative to FTP that uses SSL for encryption?

Answers:

***FTPS**

SFTP

SCP

RCP

Explanation:

FTP Secure (FTPS) adds SSL or TLS to FTP to secure login credentials and encrypt data transfers. FTPS requires a server certificate.

Secure Shell File Transfer Protocol (SFTP) is a file transfer protocol that uses Secure Shell version 2(SSHv2) to secure data transfers. SFTP is not FTP that uses SSH, but rather a secure transfer protocol that is different from FTP.

Secure Copy Protocol (SCP) uses Secure Shell version 1 (SSHv1) to secure file transfers and login credentials.

Remote Copy Protocol (RCP) is an unsecured protocol for file transfer.

q_sec_prot_ftps_02_secp8

As a network administrator, you are asked to recommend a secure method for transferring data between hosts on a network.

Which of the following protocols would you recommend? (Select two.)

Answers:

RCP

***SCP**

FTP

***SFTP**

TDP

Explanation:

The Secure File Transfer Protocol (SFTP) is a file transfer protocol that uses Secure Shell (SSH) to secure data transfers. SSH ensures that SFTP transmissions use encrypted commands and data, which prevents data from being transmitted over the network in cleartext. The Secure Copy (SCP) protocol is associated with Unix/Linux networks and is used to transfer files between systems. Like SFTP, SCP relies on SSH to ensure that data and passwords are not transmitted over the network in clear text.

The Remote Copy Protocol (RCP) and the File Transfer Protocol (FTP) are used to transfer files between computers. However, both are unsecured protocols and transmit data over the network in cleartext. Data and passwords sent over the network in clear text are in danger of being tampered with or read during transmission, making them inappropriate for many network applications.

The Tag Distribution Protocol (TDP) is a two-party protocol that runs over a connection-oriented transport layer with guaranteed sequential delivery. Tag-switching routers use this protocol to communicate tag-binding information to their peers. However, it is not a protocol used as a secure method for transferring data between hosts on a network.

q_sec_prot_https_01_secp8

Which of the following protocols uses port 443?

Answers:

***HTTPS**

S/MIME

SSH

S-HTTP

Explanation:

Hyper Text Transfer Protocol Secure (HTTPS) is a secure form of HTTP that uses SSL or TLS to encrypt sensitive data before it is transmitted. HTTPS uses port 443.

Secure Hypertext Transfer Protocol (S-HTTP) supports a wide variety of encryption methods, but it does not use port 443.

SSH uses port 22. S/MIME is a method for encrypting emails.

S/MIME does not communicate over a specific port number.

q_sec_prot_https_02_secp8

Which TCP/IP protocol is a secure form of HTTP that uses SSL as a sub-layer for security?

Answers:

***HTTPS**

SMTP

SSH

DNS

Explanation:

HTTPS is a secure form of HTTP that uses SSL as a sub-layer for security.

SMTP is used to route electronic mail through the internetwork.

SSH allows secure interactive control of remote systems.

DNS is a system that is distributed throughout the internetwork to provide address/name resolution.

q_sec_prot_https_03_secp8

Which protocol is used to browse a website securely?

Answers:

***HTTPS**

SSH

UDP

SIP

ARP

Explanation:

HTTPS is a secure form of HTTP that uses SSL to encrypt data before it is transmitted. HTTP is used by web browsers and web servers to exchange files (such as web pages) through the World Wide Web and intranet.

SSH is used for secure remote management.

UDP is a data transport control protocol that does not include error correction or detection mechanisms.

SIP is a protocol used by voice over IP (VoIP) to set up and terminate phone calls.

ARP is used by devices to find the IP address of a device with a known MAC address.

q_sec_prot_https_04_secp8

A defense contractor has tasked its local network administrator with securing communications between the organization's web server and clients to protect sensitive user information.

Which protocol should the network administrator choose to achieve this security objective?

Answers:

HTTP

Telnet

***HTTPS**

SSH

Explanation:

Hypertext Transfer Protocol Secure (HTTPS) is the secure alternative to HTTP. It uses encryption to protect transmitted data, making it the appropriate choice for securing sensitive user information, such as login credentials and data entered in form fields.

In contrast, HTTP, an insecure protocol, transmits data in clear text format, which makes it unsuitable for securing sensitive information.

Similarly, Telnet is another insecure protocol that the organization should avoid due to its lack of encryption, making it unsuitable for securing sensitive data and login information.

While Secure Shell (SSH) is a secure protocol for connecting to servers and equipment, it is not the appropriate choice for securing communication between a web server and clients.

q_sec_prot_https_05_secp8

A rapidly growing startup relies primarily on web servers for its operations, focusing almost exclusively on Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) traffic.

In the context of this focus and the need for enhanced security, which action should the company prioritize to ensure controlled outgoing network communication?

Answers:

***Allow only ports 80 and 443 for incoming traffic and restrict unnecessary outgoing protocols.**

Permit all incoming and outgoing traffic.

Restrict outgoing traffic to only port 25.

Block only incoming requests from private Internet Protocol addresses.

Explanation:

Allowing ports 80 and 443 for incoming and restricting unnecessary outgoing protocols aligns with the need for HTTP and HTTPS traffic and reduces potential security risks by restricting unnecessary outgoing protocols.

Permitting all incoming and outgoing traffic doesn't follow the principle of least privilege. It opens the network to potential threats by not regulating the type of traffic that enters or leaves, thus significantly increasing security risks.

Restricting all outgoing traffic to only port 25, used by Simple Mail Transfer Protocol for email transfer, is unnecessarily limiting. It doesn't consider the required HTTP and HTTPS traffic and could inadvertently block legitimate network communication.

Blocking only private incoming requests doesn't account for the specific incoming and outgoing traffic restrictions needed.

q_sec_prot_remote_01_secp8

Which of the following tools allow remote management of servers? (Select two.)

Answers:

***SSH**

***Telnet**

FTP

POP3

SSL

Explanation:

Both Telnet and SSH are tools for remote server management.

POP3 is for retrieving email from a remote server, and FTP is for transferring files.

Secure Socket Layer (SSL) secures messages being transmitted on the internet.

q_sec_prot_remote_02_secp8

You are a network administrator for a large corporation that has recently transitioned to a remote work model due to unforeseen circumstances. One of your tasks is to ensure secure remote access for all employees. You have decided to use a remote desktop protocol (RDP) for this purpose.

However, you are aware of the security risks associated with RDP and want to ensure that the connection is as secure as possible.

Which of the following actions would be the most effective in enhancing the security of your RDP connection?

Answers:

Disable Network Level Authentication (NLA).

Use a non-standard port for RDP.

Enable RDP on all company devices.

***Implement two-factor authentication.**

Explanation:

Implement two-factor authentication is the correct answer. Two-factor authentication adds an additional layer of security by requiring users to verify their identity using a second factor, typically a mobile device, in addition to their regular username and password. This makes it much more difficult for an attacker to gain access even if they have obtained a user's credentials.

Disabling Network Level Authentication (NLA) is an incorrect answer. NLA provides an extra layer of authentication before a session is established with the server, which helps to prevent unauthorized access.

Using a non-standard port for RDP is an incorrect answer. While this might obscure the RDP service from casual scanning, it does not provide any real security. Port scanning can still reveal the RDP service, and this method does nothing to protect against brute-force attacks or vulnerabilities in the RDP service itself.

Enabling RDP on all company devices is an incorrect answer. Enabling RDP on all devices increases the attack surface and could potentially allow an attacker to gain access to the network if a device is compromised. It's better to limit RDP access to only those devices that need it.

q_sec_prot_ssh_01_secp8

Which of the following protocols can be used to manage a network device from a remote connection securely?

Answers:

***SSH**

Telnet

SFTP

TLS

Explanation:

SSH allows for secure interactive control of remote systems. SSH is a secure and acceptable alternative to Telnet.

Telnet is a network protocol that allows a user on one computer to log into another computer that is part of the same network. However, it does not provide for secure management from a remote connection.

SFTP is a file transfer protocol that uses Secure Shell (SSH) to secure data transfers.

TLS ensures that messages transmitted on the internet are private and tamper-proof. TLS is often used to add security to other protocols.

q_sec_prot_ssh_02_secp8

SFTP uses which mechanism to provide security for authentication and data transfer?

Answers:

IPsec

Token devices

SSL

***SSH**

Explanation:

SSH File Transfer Protocol uses Secure Shell (SSH) to provide security for authentication and data transfer.

FTPS uses SSL to secure FTP traffic. You can also secure FTP traffic by establishing an IPsec tunnel between the client and the server, but IPsec is established independently of FTP in this case.

A token is a device that employs an encrypted key for which the encryption algorithm is the method of generating an encrypted password known to a network's authentication server. There are both software and hardware tokens. However, it does not provide security for data transfer.

Secure Sockets Layer (SSL) is an encryption security protocol for data. However, it does not provide for secure authentication.

q_sec_prot_ssh_03_secp8

Telnet is inherently insecure because its communications are in plaintext and easily intercepted.

Which of the following is an acceptable alternative to Telnet?

Answers:

SLIP

SHTTP

Remote Desktop

***SSH**

Explanation:

SSH (Secure Shell) is a secure and acceptable alternative to Telnet. SSH allows for secure interactive control of remote systems. SSH uses RSA public key cryptography for both connection and authentication. SSH uses the IDEA algorithm for encryption by default. But it is also able to use Blowfish and DES.

Serial Line Internet Protocol (SLIP) is the protocol TCP/IP uses when operating through a serial connection. It does not provide for remote connections.

Secure HTTP (SHTTP) encrypts data transfer but does not authentic to a client and is not a remote networking protocol.

Remote Desktop, while a remote control mechanism, is limited in use to a few versions of Windows and is not very secure.

q_sec_prot_ssh_04_secp8

What is the default encryption algorithm used by SSH (Secure Shell) to protect data traffic between a client and the controlled server?

Answers:

***IDEA**

DES

AES

Blowfish

Explanation:

SSH uses the IDEA algorithm for encryption by default. SSH (Secure Shell) is a secure and acceptable alternative to Telnet. SSH allows for secure interactive control of remote systems. SSH uses RSA public key cryptography for both connection and authentication.

SSH can use Blowfish or DES, but these are not the default methods.

SSH does not support AES (not SSH-2, at least).

q_sec_prot_ssl_01_secp8

You are purchasing a hard disk from an online retailer over the internet.

What does your browser MOST LIKELY use to ensure that others cannot see your credit card number on the internet?

Answers:

VPN

PPTP

***SSL**

IPsec

Explanation:

Your web browser uses SSL (Secure Sockets Layer) to ensure safe web transactions. URLs that begin with HTTPS:// trigger your web browser to use SSL.

A virtual private network (VPN) protects users by encrypting their data and masking their IP addresses. However, in this scenario, it would not normally be used to secure a credit card number during a transaction. SSL would be more likely to be used.

Point-to-Point Tunneling Protocol (PPTP) is a network protocol that enables the secure transfer of data from a remote client to a private enterprise server by creating a virtual private network (VPN) across TCP/IP-based data networks. Once again. the more common way of securing data during a transaction in a browser (such as a credit card number) is to use SSL.

IPSec is a set of communication rules or protocols for setting up secure connections over a network. However, in this scenario, the transaction is being made over the Internet through a web browser.

q_sec_prot_ssl_02_secp8

Which protocol does HTTPS use to offer greater security in web transactions?

Answers:

Kerberos

***SSL**

IPsec

Telnet

Explanation:

HTTPS uses Secure Sockets Layer (SSL) to offer greater security in web transactions.

Kerberos allows individuals communicating over a non-secure network to prove their identity to one another in a secure manner.

IPsec (IP security) is a suite of protocols for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet in a data stream.

Telnet is a tool for remote server management.

q_sec_prot_ssl_03_secp8

SSL (Secure Sockets Layer) operates at which layer of the OSI model?

Answers:

Application

Presentation

***Session**

Transport

Explanation:

SSL (Secure Sockets Layer) operates at the Session layer of the OSI model.

SSL operates over TCP port 443. SSL was developed by Netscape to secure internet-based client/server interactions. SSL authenticates the server to the client using public key cryptography and digital certificates, and this protocol encrypts the entire communication session between a server and a client. SSL can be used to protect web (HTTP) traffic as well as Telnet, FTP, and emails.

SSL does not operate at the Application layer, as this is where human interaction takes place, or the Presentation level, where the data is ensured to be in a usable format.

The Transport layer simply represents data being transmitted using various protocols, which is one layer below the Session layer.

q_sec_prot_ssl_04_secp8

When using SSL authentication, what does the client verify first when checking a server's identity?

Answers:

The certificate must be non-expiring and self-signed by the sysadmin.

***The current date and time must fall within the server's certificate-validity period.**

All DNS resolutions must point to the corporate intranet routers.

Master secrets are verifiable from asymmetric keys.

Explanation:

An SSL client first checks the server's certificate validity period. The authentication process stops if the current date and time fall outside of the validity period.

SSL clients verify a server's identity using the following steps:

The client checks the server's certificate validity period. The authentication process stops if the current date and time fall outside of the validity period.

The client verifies that the issuing certificate authority (CA) is on its list of trusted CAs.

The client uses the CA's public key to validate the CA's digital signature on the server certificate. If the digital signature can be verified, the client accepts the server certificate as a valid certificate issued by a trusted CA.

To protect against man-in-the-middle attacks, the client compares the actual DNS name of the server to the DNS name on the certificate.

q_sec_prot_ssl_05_secp8

You want to allow traveling users to connect to your private network through the internet. Users connect from various locations, including airports, hotels, and public access points like coffee shops and libraries. As such, you won't be able to configure the firewalls that might be controlling access to the internet in these locations.

Which of the following protocols would MOST likely be allowed through the widest number of firewalls?

Answers:

PPTP

L2TP

***SSL**

IPsec

PPPoE

Explanation:

Ports must be opened in firewalls to allow VPN protocols. For this reason, using SSL for the VPN often works through firewalls when other solutions do not because SSL uses port 443. Port 443 is often already open to allow HTTPS traffic. In addition, some NAT solutions do not work well with VPN connections.

PPTP uses port 1723, L2TP uses ports 1701 and 500, and IPsec uses UDP port 500 for the Internet Key Exchange (IKE) Protocol.

PPP over Ethernet (PPPoE) is used for connections with an always-on state, such as DSL or fiber-optic-running Ethernet. PPPoE is a modification of PPP that allows for the negotiation of additional parameters that aren't typically present on a regular Ethernet network. ISPs typically implement PPPoE to control and monitor internet access over broadband links.

q_sec_prot_ssl_cookies_secp8

A large e-commerce development team created a new web application to manage customer transactions. The application handles sensitive user data, including personal information and payment details.

The security team had concerns about potential security risks and took action to implement specific security measures to protect the application.

Which security measures should the team implement to safeguard the web application and user data? (Select two.)

Answers:

***Secure cookies**

***Secure socket layer (SSL) certificates**

Code signing

Input validation

Network monitor

Explanation:

Secure cookies play a crucial role in safeguarding sensitive user data during web transactions. When transmitted over encrypted connections, the "Secure" attribute of cookies reduces the risk of unauthorized access and data exposure, ensuring the protection of user privacy and maintaining the web application's integrity.

Secure Socket Layer (SSL) certificates are pivotal in securing web applications and user data through data encryption, integrity, and authentication.

Code signing is essential for software authentication but is not directly relevant to web applications and the security of user data during web transactions.

Input validation is also important for security, but it focuses on preventing data-related vulnerabilities and does not directly address the protection of user data during web transactions.

Network monitoring collects and reports on a variety of data from a computer network, including routers, switches, firewalls, load balancers, and even endpoints, like servers and workstations. However, it would not provide any security for the website or user data.

q_sec_prot_ssl_tls_secp8

Which of the following protocols are often added to other protocols to provide secure transmission of data? (Select two.)

Answers:

***TLS**

***SSL**

HTTPS

SMTP

SNMP

Explanation:

Both Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are protocols that are used with other protocols to add security. In addition, Secure Shell (SSH) can be used to add security when using unsecured protocols.

HTTPS is the secure form of HTTP that uses SSL.

SMTP is used for sending email.

SNMP is a network management protocol.

q_sec_prot_tls_01_secp8

Which of the following protocols can TLS use for key exchange? (Select two.)

Answers:

***Diffie-Hellman**

***RSA**

IKE

KEA

ECC

Explanation:

TLS uses Diffie-Hellman or RSA to exchange session keys.

SSL uses RSA or Key Exchange Protocol (KEA) for key exchange.

IPsec uses IKE for key exchange.

ECC (elliptic curve cryptography) is a method that can be used in key exchange.

q_sec_prot_tls_02_secp8

Which of the following statements about applying common security techniques to computing resources is correct?

Answers:

SSL is primarily for securing FTP communications.

HTTPS operates over port 80 by default.

***TLS 1.3 prevents downgrade attacks and reduces the number of handshake messages.**

A cipher suite in TLS 1.3 includes RSA signatures for symmetric bulk encryption.

Explanation:

Transport Layer Security (TLS) removes the ability to perform downgrade attacks, making it more secure. It also introduces changes to the handshake protocol, reducing the number of messages and speeding up connections.

Netscape developed Secure Sockets Layer (SSL) in the 1990s to address the lack of security in HTTP, not File Transfer Protocol (FTP).

HTTPS operates over port 443 by default.

In TLS 1.3, cipher suites no longer include key exchange and RSA signature algorithms for bulk encryption.

q_ipsec_ah_esp_01_secp8

IPsec is implemented through two separate protocols.

What are these protocols called? (Select two.)

Answers:

***AH**

***ESP**

SSL

EPS

L2TP

Explanation:

IPsec is implemented through two separate protocols, which are IP Authentication Header and IPsec Encapsulating Security Payload. IPsec AH provides authentication and non-repudiation services to verify that the sender is genuine and data has not been modified in transit. IPsec ESP provides data encryption services for the data within the packet.

SSL, EPS, and L2TP are not protocols associated with IPsec.

q_ipsec_ah_esp_02_secp8

Which of the following network layer protocols provides authentication and encryption services for IP-based network traffic?

Answers:

TCP

***IPsec**

SSL

L2TP

Explanation:

IPsec is a security implementation that provides security for all other TCP/IP-based protocols that operate above the Network layer. IPsec provides authentication through a protocol called IPsec Authentication Header (AH) and encryption services through a protocol called IPsec Encapsulating Security Payloads (ESP).

The Transmission Control Protocol (TCP) is a Transport layer connection protocol that provides data transmission services. It is not a secure protocol and relies on other measures, such as IPsec, to provide security.

Secure Sockets Layer (SSL) is an Application and Session layer protocol that is designed to secure network traffic from certain other protocols, such as HyperText Transfer Protocol (HTTP) and Post Office Protocol version 3 (POP3). It does not provide security for protocols lower in the TCP/IP protocol stack, such as TCP and UDP.

Layer 2 Tunneling Protocol (L2TP) is a protocol used to encapsulate Point-to-Point Protocol (PPP) traffic.

q_ipsec_ah_esp_03_secp8

What is the primary function of the IKE Protocol used with IPsec?

Answers:

***Create a security association between communicating partners.**

Encrypt packet contents.

Provide authentication services.

Provide both authentication and encryption.

Ensure dynamic key rotation and select initialization vectors (IVs).

Explanation:

Internet Key Exchange (IKE) Protocol is used with IPsec to create a security association between communicating partners. It controls the negotiation of encryption methods, identifies how keys are exchanged, and sets up other parameters that control communications.

Encapsulating Security Payload (ESP) provides both authentication and encryption, while Authentication Header (AH) provides authentication only.

IKE, by itself, ensures dynamic key rotation and selects initialization vectors (IVs). But this functionality is not completely utilized by IPsec.

q_ipsec_ike_01_secp8

The IT department at a software development company is setting up a VPN for employees.

The tunneling protocol the department uses must encrypt network connections at the packet level. The connection needs to support mutual authentication as well as low packet overhead.

Which setup should the team choose?

Answers:

***An IPsec tunnel with Internet Key Exchange (IKE) for mutual authentication.**

A Transport Layer Security (TLS) tunnel with extensible authentication protocol (EAP) for mutual authentication.

An IPsec tunnel with Secure Sockets Layer (SSL) for mutual authentication.

A TLS tunnel with PPTP (Point-to-Point Tunneling Protocol) for mutual authentication.

Explanation:

IPsec operates at the network layer and protects secure VPN connections at the packet level. It can also perform mutual authentication using IKE, meeting all the team's requirements.

While TLS tunnels can use EAP to perform mutual authentication, they operate at the application level and do not operate on the network layer, so they can only secure a connection between two applications.

IPsec can utilize SSL through IKEv2, but IPsec establishes mutual authentication itself and does not typically use SSL for that purpose.

PPTP is an easily breachable legacy protocol from the 1990s. Though it can use EAP-TLS for mutual authentication, it is not secure and can expose the connection to offline cracking.

q_ipsec_ike_02_secp8

A software development company has implemented an IPsec tunnel with Internet Key Exchange (IKE) for mutual authentication as part of its Virtual Private Network (VPN) setup for employees.

The company chooses this solution for its ability to encrypt network connections at the packet level and support mutual authentication with low packet overhead.

What benefits do the IT department likely aim to achieve with this setup?

Answers:

***Protection at the packet level with mutual authentication**

Secure application-level connections

Mutual authentication using SSL

Reliability of an established legacy protocol

Explanation:

Internet Protocol Security (IPsec) operates at the network layer, protecting secure VPN connections at the packet level. It also uses Internet Key Exchange (IKE) to perform mutual authentication, which aligns with the team's requirements.

The chosen setup does not likely benefit from using Transportation Layer Security (TLS) tunnels to secure application-level connections between two applications.

IPsec can utilize Secure Sockets Layer (SSL) through IKEv2, but it typically establishes mutual authentication and does not use SSL for that purpose.

Point-to-Point Tunneling Protocol (PPTP) is known to be easily breached and does not align with the team's requirements, despite older protocols that might appear tried-and-tested.

q_ipsec_remote_access_secp8

A company is setting up a secure remote access solution for its employees.

\Which of the following provides the MOST secure method for remote access?

Answers:

***A virtual private network (VPN) with Internet Protocol Security (IPSec)**

A jump server

A Layer 4 firewall

A proxy server

Explanation:

A virtual private network (VPN) with Internet Protocol Security (IPSec) provides a secure, encrypted connection for remote access, the most secure option.

While a jump server provides a controlled means of access, it does not offer the same level of security as a VPN with IPSec.

A Layer 4 firewall operates at the transport layer and provides security measures but cannot provide secure remote access. It controls access based on port numbers and protocols but cannot inspect traffic at a deeper level.

A proxy server operates at the application layer and provides some security measures but cannot provide secure remote access. It does not provide end-to-end encryption or advanced security measures for secure remote access.

q_ipsec_tls_vpn_secp8

A software company has a team of developers working remotely from different locations. They need to connect to the company's server to access and update code repositories.

The company wants to ensure that the connection is secure and that both the server and the client can authenticate each other.

Which type of VPN setup would be MOST appropriate for this scenario?

Answers:

Site-to-site VPN

Host-to-host VPN

***Transport layer security (TLS) VPN**

Point-to-Point Tunneling Protocol (PPTP) VPN

Explanation:

Transport Layer Security (TLS) VPN is the correct answer. A TLS VPN allows the client to connect to the remote access server using digital certificates. This allows for mutual authentication, where both server and client prove their identity to one another. It's the best choice for a scenario where multiple clients need to connect to a server securely and mutual authentication is required.

Site-to-site VPN is used to connect two or more networks, such as two offices in different locations. It wouldn't be the best choice for this scenario as the developers are working remotely from different locations and need to connect individually to the company's server.

Host-to-host VPN is used to secure traffic between two specific computers. While it could be used for a connection between a developer's computer and the company server, it wouldn't be the most efficient solution for a team of developers working from different locations.

Point-to-Point Tunneling Protocol (PPTP) VPN is an older VPN protocol that has been largely deprecated due to security vulnerabilities. It wouldn't be the best choice for a scenario where security is a priority.

q_ipsec_transport_mode_secp8

You are a senior network engineer tasked with implementing IPsec in her organization to enhance network security.

The organization has a hybrid network infrastructure with on-premise and cloud-based resources. The organization also has remote employees who access the network via VPN. You are considering using the Authentication Header (AH) protocol and are unsure about the mode of operation to use.

Which mode should you choose and why?

Answers:

***Transport mode because it provides sufficient security while maintaining high network performance.**

Tunnel mode because it provides an extra layer of security by encrypting the entire packet.

Transport mode because it is more compatible with the VPN used by remote employees.

Tunnel mode because it is more suitable for a hybrid network infrastructure.

Explanation:

Transport mode is the correct choice in this scenario. While it only encrypts the payload of the packet and not the entire packet, it provides sufficient security for the data being transmitted. Furthermore, because it is less resource-intensive than tunnel mode, it helps maintain high network performance, which is crucial in a hybrid network environment with potentially varying network speeds and capacities.

Tunnel mode does provide an extra layer of security by encrypting the entire packet. However, it is more resource-intensive and could potentially impact network performance. In a hybrid network environment where network performance is crucial, this could be a disadvantage.

Transport mode is not necessarily more compatible with VPNs used by remote employees. The compatibility with VPNs depends on various factors, including the type of VPN used and the specific VPN configuration.

While tunnel mode might seem more suitable for a hybrid network infrastructure because it encrypts the entire packet, it is not the best choice in this scenario due to its potential impact on network performance. In a hybrid network environment, maintaining high network performance is often a priority, making transport mode a better choice.

q_ipsec_tunnel_mode_secp8

You are a network administrator tasked with implementing IPsec to secure data transmission over the company's unprotected TCP/IP network.

You decide to use the Encapsulating Security Payload (ESP) protocol for its confidentiality feature. However, you are unsure whether to use transport mode or tunnel mode. The company's network includes several routers and switches, and it also uses Network Address Translation (NAT).

Which mode should you choose and why?

Answers:

Transport mode because it only encrypts the payload, making it faster and more efficient.

Tunnel mode because it encrypts the entire packet, providing an extra layer of security.

Transport mode because it is more compatible with NAT.

***Tunnel mode because it is more compatible with NAT.**

Explanation:

Tunnel mode is the correct choice because it is more compatible with NAT. In tunnel mode, the entire original IP packet (including the header) is encrypted and encapsulated into a new packet. This allows the packet to traverse through NAT devices without the NAT modifications causing issues, as the modifications are made to the outer packet, not the encrypted original packet.

Transport mode only encrypts the payload, not the entire IP packet. While this might be faster and more efficient, it does not provide the level of security needed when traversing multiple network devices and dealing with NAT.

Tunnel mode does provide an extra layer of security by encrypting the entire IP packet. However, the reason to choose tunnel mode in this situation is not solely because of the extra security it provides.

Transport mode is not more compatible with NAT. NAT changes IP headers, which can cause communication errors with an IPsec VPN tunnel. Transport mode, which does not encrypt the IP header, can have issues when the IP header is modified by NAT.

q_ipsec_vpn_site-to-site_secp8

A multinational corporation with headquarters in New York and branches in London, Tokyo, and Sydney wants to securely connect all their offices to share resources and data.

The IT department has been tasked with setting up a secure network connection.

Which type of VPN setup would be the MOST appropriate for this scenario?

Answers:

Remote access VPN

Host-to-host VPN

***Site-to-site VPN**

Client-to-site VPN

Explanation:

Site-to-site VPN is the correct answer. A site-to-site VPN connects two or more networks, such as two offices in different locations. In this scenario, it would allow the corporation's offices in New York, London, Tokyo, and Sydney to share resources and data securely.

Remote access VPN is typically used by individual users or telecommuters to connect to a corporate network from a remote location. While it would allow employees to access the company's network from home or while traveling, it would not connect the company's offices to each other.

Host-to-host VPN is used to secure traffic between two specific computers, typically in a situation where the private network is not trusted. It would not be suitable for connecting multiple sites of a multinational corporation.

Client-to-site VPN is similar to a remote access VPN, where individual users or devices connect to a corporate network from a remote location. It would not be suitable for connecting multiple offices of a multinational corporation.

8.8 Web Application Security

As you study this section, answer the following questions:

What are the common forms of web application attacks?

How do you mitigate replay attacks?

What are some methods to prevent driver manipulation?

How does SSL stripping work?

In this section, you will learn to:

Clear the browser cache.

Prevent cross-site scripting.

Exploit SQL on a webpage.

Perform an SQL injection attack.

The key terms for this section include:

Term	Definition
Privilege escalation	The exploitation of a misconfiguration, a bug, or design flaw to gain unauthorized access to resources.
Pointer/object dereferencing	An attack that retrieves a value stored in memory that can be exploited through a NULL pointer dereference.
Buffer overflow	An attack that exploits an operating system or an application that does not properly enforce boundaries for inputting data such as the amount of data or the type of data.
Resource exhaustion	An attack that focuses on depleting the resources of a network to create a denial-of-service to legitimate users.
Memory leak	A leak that happens when dynamic memory is allocated in a program, but no pointers are connected to it causing it to never be returned when requested.
Race conditions	A sequence of events with dependencies that a system is programmed to run in a certain order which can lead to a time-of-check to time-of-use bug vulnerability.
Error handling	The procedures in a program that respond to irregular input or conditions.
Improper input handling	The lack of validation, sanitization, filtering, decoding, or encoding of input data.
Replay attack	An attack that happens when network traffic is intercepted by an unauthorized person who then delays or replays the communication to its original receiver, acting as the original sender. The original sender is unaware of this occurrence.
Pass the hash	An attack in which an attacker obtains a hashed password and uses it to gain unauthorized access.
API attacks	A malicious use of an API (application programming interface).
SSL stripping	An attack that focuses on stripping the security from HTTPS-enabled websites.
Driver manipulation	An attack that focuses on device drivers. The attack uses refactoring or shimming.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.3 Explain various types of vulnerabilities. Application Memory injection Buffer overflow Race conditions Time-of-check (TOC) Time-of-use (TOU) Malicious update Web-based Structured Query Language injection (SQLi) Cross-site scripting (XSS) Hardware Zero-day 2.4 Given a scenario, analyze indicators of malicious activity. Application attacks Injection Buffer overflow Replay Privilege escalation Directory traversal 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Patching 3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

	Infrastructure considerations Intrusion prevention system (IPS)/intrusion detection system (IDS) 4.1 Given a scenario, apply common security techniques to computing resources. Application security Secure cookies 4.6 Given a scenario, implement and maintain identity and access management. Single sign-on (SSO) Lightweight Directory Access Protocol (LDAP)
TestOut Security Pro	3.2 Implement application defenses 3.2.3 Configure web application security 3.2.5 Configure browser settings

8.8.1 Web Application Attacks (Lesson Video)

Transcript:

Applications are used extensively in all devices, including computers, servers, mobile devices, and more. Oftentimes, these applications have access to critical components of the operating system. Attackers will target application vulnerabilities to gain elevated privileges, crash the system, access the network, steal sensitive data, and more. In this lesson, we'll look at some of the common vulnerabilities and attacks you need to be aware of and some steps you can take to mitigate them.

When an attacker gains access to a system, one of their first goals is to gain enough privileges to execute their own code on the target system, known as arbitrary code execution. There are two types of privilege escalation.
The first is horizontal escalation, which happens when the attacker accesses functionality or data that belongs to another user with the same privilege level they currently have, such as a co-worker. This type of escalation is usually only focused on gaining access to that particular account.
Vertical escalation is the more common escalation type. It occurs when the attacker gains access to an account with a higher-level access than they currently have, such as an administrator account. Vigilant monitoring and proper implementation of access control security measures, such as strong passwords and implementing the principle of least privilege, will go a long way in protecting against escalation attacks.
Now, let's look at memory injection attacks. These attacks occur when the attacker injects malicious code into a running application's process memory. One of the more common types is an overflow attack. An overflow occurs when a program is designed only to handle a certain amount of data, but the attacker sends more data, which can cause the system to crash or behave in unexpected ways. The most common of these is the buffer stack overflow, which targets the stack of a program. To better understand this, let's look at how programs are put together.
Overflow attacks are one of the most dangerous types of attacks. The main cause of these attacks is errors in programming. The best way to defend against these attacks is to use proper programming practices and avoid programming functions known to be vulnerable. Thankfully, many modern languages, such as Python and Java, have built-in controls to prevent overflow attacks.
A replay attack, also known as a session replay or session hijack, exploits the session cookies that are generated when you connect to a web application. These cookies allow the application to identify the user and prove that they've been

authenticated. The hacker can intercept the cookie using an on-path attack and then forge their own cookie file to establish a session with the web application. This allows the attacker to gain access to the victim's account without knowing their credentials.
To prevent this form of attack, use a strong digital signature with timestamps, session keys that are time-bound and process-bound, and sequence numbers. Authentication systems can be programmed to accept network packets with valid timestamps and sequence numbers.
The last attack we'll look at is the malicious update. This occurs when a software program update seems legitimate but contains malicious code. This can happen when the software's website is hacked, and the update file is replaced with the malicious one. This type of attack is particularly devastating because it not only harms the victims but can also destroy the reputation of the software developer. Malicious updates can be difficult to protect against, but secure software supply chain management, digital signature verification, and other software security practices help mitigate these risks.
That'll wrap up this lesson on web application attacks. In this lesson, we reviewed some common attacks you should be aware of, including escalation attacks, memory injection attacks, replay attacks, race conditions, and malicious updates. We also looked at some security measures we can take to help mitigate each of these attacks.

8.8.2 XSS and CSRF Attacks (Lesson Video)

Transcript:

FILE NOT FOUND

8.8.3 Injection Attacks (Lesson Video)

Transcript:

In this lesson, we're going to talk about injection attacks. We're going to focus on LDAP, XML, and command injection attacks and a variation called the directory traversal attack.

A lightweight directory access protocol, or LDAP, injection attack is used to exploit a web-based application. If the application fails to properly validate the input in the form fields, then it's possible for the wrong commands to be constructed. This could potentially provide an attacker with unauthorized access to the LDAP directory tree. They may be able to view more information than they should have access to, or even modify it.
Suppose we have an application that uses LDAP calls to communicate with the directory tree. Client workstations can access the web server, load the application's web page, and query the directory services tree.
Here, if the user supplies an invalid value, or a malicious value, unexpected behavior can occur. For example, suppose the user puts an asterisk (*) in this field and submits the form. Depending on the configuration, the directory server could respond to the LDAP query with a list of every user in the directory services tree. You do not want an attacker to be able to see all the users in your tree.
An attacker could even use this field to view the attributes of a user object, such as the password.
An XML injection attack manipulates the XML structure used by an XML-based application or service. It inserts malicious content into the structure, which can affect the application's output. An XML injection attack might cause your CDATA field to be ignored by your XML parser and input validation filters. It will slip by. In this example, the payload contains a malicious URL. If the user clicks on the link thinking it's for something else, it will actually go to that user's banking account and run a cross-site request forgery attack. This is just a simple example, and the result of the attack would depend on what the malicious URL contains.
Let's review command-injection attacks. The purpose of this attack is to inject and execute commands in a vulnerable application. The vulnerable application runs commands on behalf of the attacker with the privileges assigned to the application itself.
Let's suppose a Linux server is running the DNS service. If an attacker was able to exploit this service, they could then run commands on this Linux server. This was a very popular attack several years ago. Many system administrators ran this service as the root user account, which allowed them to run commands with full access. To combat this attack, you should never run a service as an administrative user. Instead, create a limited user account on the system that only has access to what the service requires and nothing more.

Let's look at one more variation of the injection attack called the directory traversal attack or path traversal attack. This attack's goal is to gain access to a specific file stored elsewhere in the file system. It usually exploits some type of vulnerability in the application itself, such as not properly validating user input. In this case, characters representing the path to the parent directory on the application server could be manipulated to allow access to secure files.
For example, let's say there is a file directory with two subdirectories. Inside subdirectory A, there is a file named 123. Inside subdirectory B, there is a file named 456. The user, in this case, only has access to file 456, and he can search the directory hierarchy for a vulnerability and access file 123.
Because of this vulnerability, web applications need to be coded such that they look for all characters that could be used for directory traversal, whether it's the actual characters ("..", "/", "\") or their percent encoded equivalents.
In this lesson, we looked at injection attacks. We have covered LDAP, XML, and command injection attacks. We ended by looking at directory traversal attacks.

8.8.4 Zero Day Application Attacks (Lesson Video)

Transcript:

Let's talk about the Zero-Day Attack.

A zero-day attack is a threat that exploits a vulnerability in a computer application that is still unknown, or that has never been exploited before. There are two primary vehicles used to conduct these types of attacks.
The first one is a web browser. Browsers are popular targets because almost every system has one.
The other type is sending zero-day exploits via email. A zero-day email contains a link that launches the attack if it's clicked.
A zero-day attack occurs during the window between the time when the attacker first exploits the vulnerability and when the software developer creates a fix for it. This window may vary from a couple of days to many years.
With this in mind, let's look at a process that many software developers use to prevent zero-day attacks. It is the Unknown Vulnerability Management Process
The unknown vulnerability management process is a quality assurance process that finds and fixes unknown zero-day vulnerabilities. It is composed of four phases: analyze, test, report, and mitigate.
During the Analyze phase, the developers are trying to find the application's vulnerabilities. They're looking for weaknesses and attack vectors, which are ways an attacker can get into the application. Here, the developers try to break the application themselves.
Next is the Test phase. During this phase, what's called fuzz testing conducted against the attack vectors identified during the analysis phase.
After testing comes the Report phase. During the report phase, personnel try to reproduce any issues found during the testing phase to make sure that the situation is replicable and represents a real vulnerability.
Next is the Mitigate phase. This is where a fix is developed for the vulnerability, based on the results of the testing and reporting phases.
What should Security Administrators do to defend against zero-day attacks? First is to configure automatic updates for the operating system and the application.
By configuring automatic updates to the operating system, you are patching holes in the operating system where the application is running. The zero-day attacker is going to be looking for an exploit unique to that application. However, you want to close as many attack vectors as you can. Keeping your operating system updated will make your system safer from attacks in general.
You should also have a network firewall and a host firewall. By using a host-based firewall, you're hardening the system where the application is running, making it more difficult for an exploit to take place.
Finally, you should use an IDS or an IPS. For zero-day attacks, you are looking for anomalies. Consequently, virus scanners or malware detection applications probably won't work because they're based on signatures. There won't be a signature for an unknown exploit, and the attack will slip by your malware detection software.
By using an intrusion detection system or an intrusion prevention system, you look for anomalies against the norm. The IDS system will alert you that there is something unusual going on with the system. You can then begin troubleshooting, and you may discover a potential zero-day attack occurring.
In this video, we covered the basics of zero-day attacks and what they are. Then we looked at the methodology and procedures that software developers use to find the unknown vulnerabilities. Finally, we looked at a few things that you can do on your end to defend against zero-day attacks.

8.8.5 Web Browser Security Facts

This lesson covers the following topics:

- Manage browser data
- Enhance browser privacy

Manage Browser Data

A *web browser* is an application for retrieving and displaying information on the internet. Web browsers present the possibility of security breaches in an organization's network. There are general actions and browser-specific actions you can take to help harden the network against attacks from the internet.

When using a browser, the following might indicate an unsecured connection or an attack.

- A web document with a URL that contains a new or different domain name than the site you intended to visit.
- A menu bar that includes new commands or is missing common commands.
- The status line of the browser displays an unlocked symbol when SSL should be in use.

Regardless of the browser you are using, clear your private data regularly. Private data can be cleared based on the data's age. You can clear data from the last few hours, the last few weeks, or all the time. The type of browser you are using, and the types of sites you have visited will determine the type of data that can be cleared. Most browsers let you clear the following data:

- Browsing history
- Download history
- Cookies and other site data
- Cached images and files
- Passwords
- Auto-fill form data
- Site permissions
- Hosted app data

The following table lists steps for each browser to clear data.

Browser	Steps
Google Chrome	Steps for the Google Chrome browser are:

	Select the ellipses (three dots) button on the menu bar. Go to History > History. Select **Clear browsing data** .
Microsoft Edge	Steps for the Microsoft Edge browser are: Select the ellipses (three dots) button on the menu bar. Go to History. Select **Clear Browsing Data.**
Internet Explorer	Steps for the Internet Explorer browser are Select the Tools (gear) icon from the menu bar. Select Internet Options. Go to Browsing history. Select **Delete** .

Enhance Browser Privacy

You can use the following browser settings and guidelines to enhance browsing privacy and security. These may be named and implemented differently in different browsers, but the general ideas are the same.

Settings	Description
Cookies	*Cookies* are text files that save information about preferences, browser settings, and web page preferences. They identify you (or your browser) to websites. Be aware of the following facts about cookies: Cookies aren't inherently malicious and are often necessary for e-commerce websites. The use of cookies can constitute a privacy violation because cookies can retain personal information. A hacker could gain access to this information. Cookies can be misused by malware to collect and report your web surfing activities. *First-party cookies* are cookies used by the site you are visiting.

	Third-party cookies are cookies placed by sites linked to the site you are visiting. For example, banner ads on a website might place cookies on the machine to identify ads already seen or ads opened. Secured environments should restrict the use of cookies on all web browsers and other internet service utilities. Cookies can usually be found in the user profile in the file system.
Cache	A *cache* is a storage location for information that will be used again, such as images, sounds, web pages, and even usernames and passwords used on websites. In addition to taking up space, data in the cache could be retrieved by someone with access to the computer. To provide some level of protection, you should clear the web browser cache whenever you use a public computer to access the internet, especially when you have accessed sites for retrieving personal data.
Security	Enable the following options to increase security: Warn me when sites try to install add-ons. Block reported attack sites. Block reported web forgeries. It is best practice always to enter passwords and not to have the browser remember them. Do not select the **Remember passwords for sites** option. Do not select the **Use a master password** option. When you select this option, all passwords saved on the system are encrypted. You create a master password that retrieves and unencrypts passwords for individual sites.
Add-ons	An *add-on* , also known as a *plug-in* or *browser extension* , is a program that adds functionality and features to a web browser, including extra toolbars and interactive web content. Over time, a browser collects add-ons, some of which could have malicious intent. Secure the browser by reviewing add-ons and uninstalling items that are not appropriate for the environment. *Disabling* an add-on disables it for the current user. This allows users to enable or disable add-ons based on their own needs. *Deleting* an add-on removes it from the system and prevents any user from using it.
General	General information for web browser security includes: Use the **Always ask me where to save files** option to avoid having files downloaded without your knowledge. By using this option, you will always know when a file is being downloaded to the system. Enable the **Block Pop-up windows** option. Turn off **Remember search and form history** . Data you enter into forms, such as your banking account number, will be stored if this option is on.

	Turn off **Accept third-party cookies** or accept cookies and specify **ask me every time** so you will know when third-party cookies are created.

8.8.6 Clear the Browser Cache (Simulation)

Scenario

You use Google Chrome as your web browser on the desktop computer in your dorm room. You are concerned about privacy and security while surfing the web. You are also concerned about exploits that harvest data from your Google Chrome browsing history.

In this lab, your task is to delete the following items from your Google Chrome browser history for all time:

- Browsing history
- Download history
- Cookies and other site data
- Cached images and files
- Hosted app data

Explanation

Complete this lab as follows:

1. Delete all items from your Google Chrome history.
   1. From the Windows taskbar, select **Google Chrome** .
   2. In the upper right, select the **ellipsis** (three dots) and then select **History** > **History** .
   3. Maximize the window for better viewing.
   4. Select **Clear browsing data** .
   5. Select **Advanced** .
   6. For the Time range field, use the drop-down menu to select **All time** .
   7. Make sure the following items are checked:
      - Browsing history
      - Download history
      - Cookies and other site data

Cached images and files

Hosted app data

Select **Clear data** .

8.8.7 Preventing Cross-Site Scripting (Demo Video)

Transcript:

In this demonstration, we'll explore the steps you can take to safeguard yourself against cross-site scripting attacks, commonly known as XSS attacks. Although XSS attacks can exploit various web technologies like HTML and Flash, JavaScript is one of their primary vehicles. Consequently, we'll focus on how to disable JavaScript in Microsoft Edge and Google Chrome.

Before we delve deeper into this demo, please be aware that you can indeed prevent cross-site scripting attacks by disabling JavaScript in web browsers such as Microsoft Edge and Google Chrome. However, it's crucial to understand that disabling JavaScript can significantly impact the functionality and user experience of many websites. Please consider this trade-off carefully before implementing such restrictions.
Let's begin with Microsoft Edge. I've already opened Edge, and now we'll navigate to the settings. Click on the three dots, also known as the ellipsis, and select Settings.
In the settings menu, go to Cookies and site permissions. You'll find a list of items on the right. Scroll down until you locate JavaScript. Here, you can toggle off the switch next to Allowed (recommended) to disable JavaScript. In some versions of Edge, you might receive a confirmation dialog; if so, click Turn off to confirm the action. JavaScript is now disabled in Microsoft Edge. To enable it again, simply follow the same steps and toggle the switch back on.
Let's switch to Google Chrome and see how to disable JavaScript there. Click on the three dots in the top right corner, then select Settings. Navigate to Privacy and security and click on Site settings. Scroll down until you find JavaScript, and click on it. Choose the radio button that says, Don't allow sites to use JavaScript to disable it. To enable JavaScript, simply select Sites can use JavaScript to turn it back on.
It's also essential to mention the potential risks associated with the browser's credential autofill feature. This feature is often enabled by default in most commonly used browsers. To disable the ability to save passwords, go to Autofill and passwords and click on it. If you use Google Password Manager, you can toggle off Google's offer to save passwords and the ability to sign in automatically. Remember that these settings can be adjusted as needed.
In summary, the most effective prevention of cross-site scripting is typically implemented in the design of web applications. However, as a user, you may not always know if the sites you visit have taken all necessary precautions to protect against XSS attacks. Therefore, being able to disable JavaScript in your browser can provide an additional layer of security when needed. Just remember to weigh the potential loss of website functionality against the security benefits when deciding to disable JavaScript.

8.8.8 SQL Injections (Lesson Video)

Transcript:

SQL injection attacks are powerful and complex. They're the tool behind many successful high-profile internet security breaches.

In 2011, Sony was targeted by an SQL injection attack that compromised over a million emails, usernames, and passwords. In 2013, the United Stated Department of Energy was targeted, and at least 100,000 employee records were compromised. These records included contact information, social security numbers, and even bank account numbers. Perhaps the most ironic attack occurred when mysql.com was attacked, releasing a large list of usernames and passwords to hackers. In each of these instances, the vulnerability wasn't in the SQL software, but in the way the websites and applications were implemented. Each of these attacks could have been prevented with careful front-end configuration and penetration testing.

Because the SQL injection attacks target web applications, let's review how they work. When a user connects to a web application, they make a request through the browser. This request travels over the internet and to the web server. The web server accepts the request and sends it to the corresponding web app.
The web application accesses the database, completes the requested task, and then responds to the web server. Once the transaction is complete, the web server sends the requested information to the user's browser.
Because attackers are seeking information stored in the database, we also need to review how those work. Databases store all sorts of information, including application data, configuration data, customer data, login information, and--well--the possibilities are endless, so you can imagine why databases are valuable targets.
A database is typically described by the way it stores data. A relational database can be organized in different ways, depending on need. For example, customer orders can be sorted by customer number, zip code, or product number. A distributed database is designed so it's easy to replicate to various locations across a network. An object-oriented database is designed around object classes and sub-classes.
Inside these databases are various methods used to organize, manage, and retrieve data. Records and rows are used to represent a collection of relative data, such as information about a product, a user, or a customer. Think about the times you've gone shopping online, and you've sorted your selection by color, cost, or rating. These filters are possible because of these databases. When you think about the number of times you've entered contact information, account information, or other personal information, you can start to understand how valuable these stores of data are.
SQL was specifically designed to request data from a database. These requests take the form of a query--a question--that asks the database to provide information specific to your request.
It's important to note that SQL injections are a result of flaws in web applications, not in the database or the web server. These attacks exploit non-validated input vulnerabilities and use them to send SQL commands through the web application and to the database. This is done by injecting a code into an existing line of code before sending it onto the database for execution. If the injection is successful, the malicious code runs on the backend database and returns the requested information.
So, how does an SQL injection work? Let's start with a very basic example.
Let's say you're logging into your account. Normally, the user will enter their username-bobsmith--and their password-secret--into the appropriate spaces. Once they click Submit, the web application will send a string to the web server that contains the credentials.
The command tells the database to check for the provided username and compare it to the stored password before granting access. Assuming that the user is found, they're directed to the requested page. As you can see, the data that the user entered is put into the same query as the commands. As a result, this code is susceptible to SQL injection attacks. If the login fields haven't been restricted, an attacker can add anything they want to them.
Let's say the attacker knows that Bob's username is bobsmith, but they don't know Bob's password. They may be able to enter 'bobsmith'--'in the username field. The single quote indicates that data has ended and a command is beginning. The double dashes indicate that code is ending and a comment is being entered. Comments are code that a program doesn't execute--they usually contain explanations or reminders for the programmer. Because of this, the application knows to ignore the comments.
Now, because the command treats everything after the dashes as comments, the instruction to verify the username with a given password is no longer visible, and the user is granted access to the user account.
Although it's fairly easy to detect the initial vulnerabilities for this attack, SQL injections are extremely complex. They require a lot of patience and a high level of database and coding expertise. So, why would attackers go to all this trouble? Well, with great effort comes great reward, and that's what these attackers are counting on. SQL injections can be used to implement several types of attacks.
When an attacker attempts an authentication bypass, they log into an application with administrative privileges without having to give a valid username or password. An information disclosure attack provides an attacker with sensitive information from the database. An attacker can use a compromised data integrity attack to deface a web page or alter a database's contents. Attackers can use an SQL injection attack to compromise data's availability; in other words, they can use it to delete information stored in the database. Attackers can also use this type of attack to remotely execute code and compromise the host operating system.
Well, that's all for this lesson. We talked about SQL injections and how they relate to web application and database technologies.

8.8.9 Exploit SQL on a Web Page (Demo Video)

Transcript:

The internet used to be pretty simple and straightforward. There was no JavaScript, Flash, CSS, backend databases, or, really, any complex web design technologies. That's all changed, and it's changing more rapidly than ever. These days, most websites are dynamic and database-driven. Site content is dependent on user input most of the time, and that data is written to a database, typically an SQL database. Like every other piece of computing technology that's invented to solve a problem, hackers have figured out how to attack these databases.

The term for attacking an SQL database is an SQL injection. It's a very common way to attack databases. The Open Web Application Security Project, OWASP, almost always includes weak SQL databases in their list of the top ten most widely exploited vulnerabilities.
In this demo, we're going to cover a few SQL injections on a vulnerable web page. This isn't a demo on SQL in general. We're just taking a quick look at some things to be aware of as a penetration tester or ethical hacker.
I'm on a Windows 10 system. I've already done a few things to it. First, I've downloaded and installed XAMPP, which is a package that contains Apache, Maria Data Base, PHP, and Perl. I've also downloaded and configured the Darn Vulnerable Web Application, DVWA, which is a vulnerable PHP and MySQL web application designed for security professionals to practice their penetration testing and ethical hacking skills.
First, we want to do a few simple injections. Let's start out by typing in a number, number 1. Click Submit. Now let's type in 2 and click Submit. All this does is give you the user ID, first name, and surname. That's not very useful, but it does show that it's not designed very well. We want to get more than that, so let's expand on our SQL injection.
We're going to enter code to pull up records for a specific user. By crafting our input, we can get more from the output than the code author intended. So, let's enter in the number 1, which is the user ID, with a single quote followed by a true statement. Our true statement can be anything that's true, such as 1 is equal to 1. It will look like this: â€˜1' and 1=1#'. Click Submit. We see the first name and the surname for that user.
Now let's go a little further. Let's say we want to find out our database name and username for the database. For that, I'll do a select statement by entering in â€˜1' and 1=1 union select database(), user()#'and press Enter. Now, down here, I get the name of my database, which is dvwa, and the username is root@localhost.
I want to see if I can get a list of tables from our database. For that, I'll do another select statement and enter in â€˜1' and 1= union select null, table_name from information_schema.tables#' and press Submit. Okay. I just got a long list of tables here. I'm looking for something that might have a list of usernames and passwords in it. Right here, I see a table called users, so I wonder if that might be the list of users along with their passwords? There's one way to find out: let's see if we can pull information from that table.
To pull the information out of the Users table, I'll enter in the following select statement: â€˜1' and 1= union select user password from users#' and click Submit. Here, you can see a nice list of the usernames along with the password hashes for each one. Now we can use a cracking program to attempt to crack any of these hashes.
And that's it for this demo. In this demo, we used a vulnerable website to practice some SQL injections. First, we found a list of users. Then we expanded to find the name of the database and the username. We ended by finding a list of tables, specifically the user table, and then retrieved the usernames and password hashes from that table.

8.8.10 Web Application Attack Facts

People today connect, learn, shop, provide services and information, and do business over the internet. All of this is made possible through web browsers and web applications. There are literally thousands of applications that are used in our everyday lives. With so many options, there are many ways that attackers have found to exploit them.

This lesson covers the following topics:

- Privilege escalation
- Pointer/object dereference
- Buffer overflows
- Resource exhaustion
- Memory leak

Race conditions

Error handling

Improper input handling

Replay attacks

Pass the hash

Application programming interface (API) attacks

Secure Sockets Layer (SSL) stripping

Driver manipulation

Privilege Escalation

Most attacks are some form of privilege escalation. There are two types:

Horizontal

This is when an attacker gains data that belongs to another user with the same privilege level as themselves (like a co-worker).

Vertical

This is when an attacker uses a system's vulnerabilities to escalate privileges to gain administrative access.

Pointer/Object Dereference

Dereferencing a pointer is retrieving the value stored in memory.

Some important facts:

A pointer stores a memory address.

All operating systems embed the kernel in the user's workspace.

The kernel is the operating system's core program that controls everything in the system.

Page protections protect the kernel from user access but can be exploited by a DoS attack through a NULL pointer dereference.

If a DMA driver module does not have enough security protections in place, it can release user pages that are pinned to a pointer with a NULL value. This happens when:

An app dereferences an object that comes back NULL instead of valid.

Null is exploited as a constant built in to evaluate to 0 in the C language.

An x86 system has a valid 0 address in the kernel address space.

Buffer Overflows

Buffer overflow important facts:

- A buffer is a temporary data storage area with limited space.
- Overflows occur when more data is attempting to be stored than the program was written for.
- Can allow hackers to cause data to flow to other memory areas that may not be protected.
  - Attackers may now access database files or system files and can replace executable code with malicious code. This is called arbitrary code execution.
- It can cause DoS attacks by crashing the program.
  - IT can occur in routers, IoT devices, and firewalls.

Resource Exhaustion

Resource exhaustion is a form of attack that focuses on depleting the resources of a network to create a denial of service to legitimate users.

This attack can be done through:

- Slow header attacks:
  - Send HTTP headers so slowly that it prevents other users from accessing the site.
  - Can be prevented with HTTP header timeouts.
- Slow post attacks:
  - Send HTTP POST body very slowly. This is done through forms, logins, and feedback input fields.
  - Can be prevented by setting a maximum body size for each form and setting the web server setting with a max total transfer time.

Resource exhaustion attacks can be focused on memory, file system storage, database connection pool entries, or the CPU. When an allocation of these resources is requested, but the size of the resource or number is not controlled, a denial of service results from a lack of resources.

Memory Leak

A memory leak happens when dynamic memory is allocated in a program, but no pointers are connected to it. This causes it never to be returned when requested.

- Programmers often create temporary memory allocations. This becomes a problem when they are not deleted after use.

Whether unintentionally leftover from a project or intentionally created by an attacker, memory leaks can result in:

- Resource exhaustion.
- DoS.
- Exploitation of other areas affected by low-memory conditions.

To mitigate these attacks:

- Delete unneeded memory allocations when finished with a project.
- Ensure that pointers are properly connected to memory values.

Race Conditions

Another web application vulnerability is a time-of-check to time-of-use bug, or (TOCTTOU) bug. This happens when a system is programmed to run with certain processes dependent on a sequence of events or race conditions.

It can happen when an attacker schedules an execution of operation between a time of check and a time of use and forces the user's process to pause or send an error. For example, in the moment between authenticating to a system and utilizing the system, the attacker can jump into the process and act as the authenticated user, leading to privilege escalation.

To mitigate:

- Ensure your operating system's file system state is not allowed to change between two system calls.
- Use file system calls that run on file handles instead of file names when possible.
- Lock single files before the check.

Error Handling

Improper error handling can create vulnerabilities in a system by revealing information that attackers can use to exploit the system. This display of too much information can result from coding practices that are not in alignment with security policies. Some examples are:

- An attacker may use a SQL injection attack that fails initially. However, the error message discloses the malformed query, which could show the query logic or other sensitive data, like passwords. The attacker can use the new information from the error message to gain access to the system.
- The disclosure of the full pathname in an error message generated from a path-traversal weakness exploit attempt.

To mitigate, be sure to program the error message with minimal information that's only useful to the intended audience.

Improper Input Handling

Improper input handling refers to a lack of validation, sanitization, filtering, decoding input data, or encoding input data.

Processing of untrustworthy input data can lead to:

- Buffer overflows.
- XSS.
- Directory traversal.
- NULL byte injections.
- SQL injection.
- Uncontrolled format string.
- DoS.
- OS commanding.

To mitigate:

- Set specific parameters for acceptable data forms and types.
- Accurately define data restrictions.
- Sanitizing, validating, and filtering properly.

Replay Attack

Replay attacks happen when network traffic is intercepted by an unauthorized person who then delays or replays the communication to its original receiver, acting as the original sender. The original sender is unaware of this occurrence.

- Also known as session replay attacks
- They are a type of man-in-the-middle attack

To mitigate, implement:

- Strong digital signatures with timestamps.
- Session keys that are time-bound and process-bound.
- Sequence numbers.
- Program authentication systems to accept network packets that have valid timestamps and sequence numbers.

Pass the Hash

Pass the hash is dangerous to an organization because once an attacker gains access, the whole organization can be compromised very quickly.

How it works:

An attacker gains access to an individual computer through malware or other techniques.

The attacker accesses the memory in the workstation to find stored hashes of other users that have used the workstation.

The attacker uses the stored hashes to gain access to other workstations in search of a station that grants privilege escalation.

To mitigate:

Use direct networking to prevent standard users and local admin users from having access to other user's workstations.

Use Group Policy Object (GPO) Editor to disable Remote Desktop Connections in an Active Directory network.

Limit domain admins' access to only workstations with the same level of privileges.

Create separate standard user-level accounts for admins when accessing lower-level privilege machines.

Application Programming Interface (API) Attacks

Application programming interfaces (APIs) are how businesses transfer information between systems within their organization or how a business communicates information to another organization. This is also a means of information transfer between companies and their customers, as APIs are the way most applications communicate with websites.

Many APIs are openly published to promote customer usage and make interactions easy. However, they do create an opportunity for a malicious user to exploit the interface to gain access to internal data and infrastructure. For example, an e-commerce site may use its API for product catalog pages on its website, in its mobile app, for a third-party reseller, and for search engine bots that bring customers to its website.

To mitigate potential API problems:

Implement rate limiting. This limits the number of calls from a client within a time limit.

Use security logs to detect and analyze unauthorized access attempts.

Look for SQL injections. These happen when a SQL statement is entered in a data field and gets executed in the database.

Make sure that program notifications are sent when there is an excess of error messages.

Secure Sockets Layer (SSL) Stripping

SSL stripping is an attack that focuses on stripping the security from HTTPS-enabled websites. This is how it works:

An attacker intercepts the initial request a user sends to a website.

The attacker establishes a secure connection with the intended server and an insecure HTTP connection with the user where all communication goes through them.

The attacker can intercept the initial request when it comes through a 302 redirect or through a non-SSL site that provides a link to a proxy that looks like the intended site.

To mitigate:

Encrypt all elements of your site with an SSL certificate.

Add your domain to the HSTS preload list. This lets browsers know that your site is secure.

Driver Manipulation

A device driver is a small piece of software that provides an interface between the operating system and a hardware device, such as a printer, keyboard, or network card. Attackers can manipulate a driver by adding malicious logic. Driver manipulation attacks often happen as a result of a web application attack, such as a drive-by download or through social engineering or phishing. The goal is to replace a good driver with one that is malevolent or to add software that comes between a good driver and the operating system.

Common driver manipulation attacks include:

Attack	Description
Refactoring	Software or code refactoring is usually considered a beneficial practice. The external behavior of refactored software code does not change. Internally, the code is modified to improve readability, reduce complexity, or improve efficiency. Attackers refactor device drivers so that their external behavior does not change. The printer, keyboard, network card, or hardware controlled by the driver still functions properly. This makes it hard to detect any problems. Internally, the refactored driver now has hidden functions that benefit the attacker.
Shimming	Like refactoring, shimming is usually beneficial. As operating systems and other software libraries are updated, their application programming interface (API) may change. The API specifies how other programs should interact with the software library or operating system. If the API is updated with new specifications, other programs using older API specifications may not work. To remedy this, a shim can be used. A shim is software placed between the newer API and software that conforms to the older API. The shim intercepts calls to the older API, translates them, and passes them to the newer API. In some cases, they can redirect the API calls elsewhere to complete the expected operation called for in the older API. Attackers can modify existing shims by injecting malicious code. They can also create a shim that intercepts valid API calls. However, the shim executes malicious code before it passes the valid calls through to the API.

To mitigate:

Use the latest browser version and patch level.

Verify that the operating system is at the latest patch level.

Install antivirus, anti-spyware, pop-up blocking, and firewall software.

Use input validation when programming services.

Client-side validation should first be used on the local system to identify input errors before the data is ever sent to the server.

For example, if the user enters an invalid value in an email address field, the error can be detected before the data is submitted.

Server-side validation should be used for error detection after the data is sent to the server. Experienced attackers can circumvent client-side validation techniques to send malicious information to the server.

For example, an attacker could send data to the server from outside the application's standard user interface, bypassing any input validation measures that may have been implemented on the client. It is unwise to rely solely on client-side input validation techniques.

Implement DNS Security Extensions, or DNSSEC. This is a security measure that only allows connection to your computer from servers that have previously been given a digital certificate.

Use HTTPS. This transfer protocol encrypts the HTTP over Transport Layer Security (TLS) or over Secure Socket Layer (SSL), protecting your browser against threats.

Use add-ons to increase the security of browsing activities:

NoScript blocks all active content except from sites you trust.

Adblock Plus blocks advertisements and ad banners (which could contain malicious code) on the internet.

Train users to log out of websites when finished. Users should never allow applications to remember their authentication information.

8.8.11 Perform an SQL Injection Attack (Simulation)

Scenario

Your name is Blake Jackson. You are the penetration tester for a small corporate network. After performing several SQL injection attract tests on the corporate network, you have decided to see how secure your own online bank's web page is.

In this lab, your task is to perform a simple SQL injection attack using the following information:

Your bank's URL: **MySecureOnlineBank.com**

Make an account query using your account number: **90342**

Answer Question 1.

Perform a simple SQL attack using: **0 OR 1=1**

Use the entire statement of " **0 OR 1=1** " but without quotes.

Answer Question 2.

Explanation

Complete this lab as follows:

Look up Blake Jackson's account balance.

From the taskbar, select **Google Chrome** .

Maximize the window for better viewing.

In the URL field, type **mysecureonlinebank.com** and then press **Enter** .

In the *Enter your Account Number* field, enter **90342** .

Select **Lookup** .

From the top right, select **Answer Questions** .

Answer Question 1.

Minimize the **Answer Questions** window.

Perform a simple SQL attack.

In the *Enter your Account Number* field, enter **0 OR 1=1** for the SQL injection.

Select **Lookup** .

From the top right, select **Answer Questions** .

Answer Question 2.

Select **Score Lab** .

8.8.12 Practice Questions (Section Quiz)

q_browser_sec_add-on_secp8

You are a cybersecurity consultant hired by a company that has recently experienced a data breach.

After an initial investigation, you discover that the breach originated from a compromised workstation where the user frequently used a web browser with multiple add-ons.

What is the most effective action to take to enhance browser privacy and prevent future breaches?

Answers:

Clear the web browser cache

Disable all browser add-ons

***Review and uninstall inappropriate add-ons**

Turn off Remember search and form history

Explanation:

Reviewing and uninstalling inappropriate add-ons is the correct answer. By reviewing and uninstalling inappropriate or potentially malicious add-ons, you can directly address the source of the problem and enhance the browser's security and privacy.

While clearing the web browser cache can help to remove stored information such as usernames, passwords, and web pages that have been visited, it does not address the issue of potentially malicious add-ons that could compromise the browser's security.

Disabling all browser add-ons can prevent potentially malicious add-ons from running, but it may also disable useful and safe add-ons. This approach is too broad and does not specifically target the source of the problem.

While turning off *Remember search and form history* can help to prevent the browser from storing data entered into forms, it does not address the issue of potentially malicious add-ons that could compromise the browser's security.

q_browser_sec_clear_data_secp8

As a cybersecurity expert, you are tasked with advising a company on best practices for managing browser data to maintain security and privacy.

Which of the following is the MOST crucial step to take?

Answers:

Regularly update the browser.

Use incognito mode for all browsing.

***Clear your private data regularly.**

Only visit HTTPS websites.

Explanation:

Clearing your private data regularly is the correct answer. Regularly clearing private data, such as browsing history, cookies, cached images, and files, can help prevent unauthorized access to this information and protect user privacy.

While keeping the browser updated is important to ensure that the latest security patches are applied, it does not directly manage the data the browser stores during use. Therefore, it is not the most crucial step.

Incognito mode prevents the browser from storing browsing history, cookies, and form data. However, it does not prevent third-party websites from tracking user activity during the session. Moreover, it does not manage data from non-incognito browsing sessions.

While visiting only HTTPS websites can help protect data during transmission by encrypting it, it does not manage the data the browser stores during use. Therefore, it is not the most crucial step.

q_browser_sec_clear_web_cache_secp8

You are the IT manager at a large corporation. One of your employees has been using a public computer to access the company's internal systems.

You notice that the employee has not been following best practices for enhancing browser privacy.

Which of the following actions would be the most effective in preventing unauthorized access to sensitive company data?

Answers:

Disabling all browser add-ons

Blocking all pop-up windows

***Clearing the web browser cache**

Turning off Remember search and form history

Explanation:

Clearing the web browser cache is the correct answer. The web browser cache can store information such as usernames, passwords, and web pages that have been visited. If this data is not cleared, someone with access to the computer could potentially retrieve this information and gain unauthorized access to sensitive company data.

While disabling all browser add-ons can help to enhance browser security by preventing potentially malicious add-ons from running, it does not directly prevent unauthorized access to sensitive company data that may have been stored in the browser cache.

Blocking all pop-up windows can help to prevent unwanted or potentially harmful pop-ups from appearing, but it does not directly prevent unauthorized access to sensitive company data that may have been stored in the browser cache.

While turning off *Remember search and form history* can help to prevent the browser from storing data entered into forms (such as banking account numbers), it does not directly prevent unauthorized access to sensitive company data that may have been stored in the browser cache.

q_browser_sec_files_download_secp8

As a cybersecurity expert, you are advising a company on best practices for general web browser security.

\Which of the following measures is MOST crucial to prevent unauthorized file downloads?

Answers:

Enable the Block Pop-up windows option.

Turn off Remember search and form history.

***Use the Always ask me where to save files option.**

Turn off Accept third-party cookies.

Explanation:

While blocking pop-up windows can help prevent unwanted or potentially harmful pop-ups from appearing, it does not directly prevent unauthorized file downloads.

While turning off *Remember search and form history* can help to prevent the browser from storing data entered into forms, it does not directly prevent unauthorized file downloads.

Using the *Always ask me where to save files* option is the correct answer. By using this option, users will always be aware when a file is being downloaded to the system, which can help prevent unauthorized or unexpected file downloads.

While turning off *Accept third-party cookies* can help to prevent third-party websites from tracking user activity, it does not directly prevent unauthorized file downloads.

q_browser_sec_web_cookies_secp8

You are a cybersecurity analyst at a large corporation. One of your responsibilities is to ensure that employees are following best practices for enhancing browser privacy.

You notice that one department is experiencing a higher than usual number of security incidents related to their web browsing habits.

What is the MOST effective action to take to enhance browser privacy and reduce these incidents?

Answers:

Enforce the use of incognito mode.

***Restrict the use of cookies.**

Clear the web browser cache regularly.

Disable all browser add-ons.

Explanation:

Restricting the use of cookies is the correct answer. Cookies are text files that save information about preferences, browser settings, and web page preferences. They can be used by websites to track user activity and can potentially be misused by hackers to collect and report your web surfing activities. Restricting the use of cookies can significantly enhance browser privacy and reduce security incidents.

While using incognito mode can help to prevent the browser from storing browsing history, cookies, and form data, it does not prevent third-party websites from tracking user activity during the session. Moreover, it does not manage data from non-incognito browsing sessions.

While clearing the web browser cache can help to remove stored information such as usernames, passwords, and web pages that have been visited, it does not address the issue of cookies, which can track user activity and potentially be misused by hackers.

Disabling all browser add-ons can prevent potentially malicious add-ons from running, but it may also disable useful and safe add-ons. This approach does not specifically target the issue of cookies, which can track user activity and potentially be misused by hackers.

q_webattk_app_attack_secp8

You are a security analyst at a tech company. You notice a significant increase in the number of failed login attempts on the company's web application. You suspect that the application might be under an application attack.

Which of the following actions should you take next?

Answers:

Ignore the failed login attempts as they are likely due to users forgetting their passwords.

***Implement a CAPTCHA system to prevent automated login attempts.**

Advise users to change their passwords immediately.

Shut down the web application until the issue is resolved.

Explanation:

Implementing a CAPTCHA system can help prevent automated login attempts, which are common in application attacks, and is the correct answer. This can help protect the application from being compromised.

Ignoring the failed login attempts could potentially allow an attacker to compromise the security of the website or web application. It's important to investigate unusual activity to ensure the security of the application.

While advising users to change their passwords can be a good security measure, it may not prevent an application attack. If the attack is automated, changing passwords will not stop the attack.

While shutting down the web application could stop the attack, it is not a practical solution. It would disrupt service for all users and may not address the underlying security issue.

q_webattk_buffer_01_secp8

Which of the following attacks is a form of software exploitation that transmits or submits a longer stream of data than the input variable is designed to handle?

Answers:

Time-of-check to time-of-use attack

Data diddling

Smurf attack

***Buffer overflow attack**

Explanation:

A buffer overflow occurs when software code receives more input than it was designed to handle. This normally occurs because the programmer of that code failed to include input validation checks. When a buffer overflow occurs, the extra data is pushed into the execution stack and processed with the security context of the system itself. In other words, a buffer overflow attack often allows the attacker to perform any operation on a system.

A time-of-check to time-of-use (TOCTOU) attack occurs when the results of an attack are realized or initiated after the attack itself is perpetrated.

Data diddling is the purposeful altering of data.

A smurf attack is a form of distributed-reflective denial of service.

q_webattk_buffer_02_secp8

A programmer that fails to check the length of input before processing leaves his code vulnerable to which form of common attack?

Answers:

Backdoor

Session hijacking

***Buffer overflow attack**

Privilege escalation

Explanation:

Buffer overflow attacks are made possible by the oversight of programmers. A simple check on the length (and sometimes format) of input data before processing eliminates buffer overflow attacks.

A backdoor is a developer-planted or cracker-planted entry device that bypasses security to gain access to a system or software. A developer-planted backdoor is often a debugging tool mistakenly left in place when the software went to market. A cracker-planted device is often a remote access server that listens for inbound connections on a specific port. Either method can be used by an intruder to gain entry into a secured environment.

Session hijacking is the concept of being able to take over a communication session between a client and server. This usually involves taking over the identity of the client and fooling the server into communicating with the pseudo-client.

Privilege escalation is the act of a user stealing or obtaining higher-level privileges in a computer system.

q_webattk_buffer_03_secp8

Having poor software development practices and failing to program input validation checks during development of custom software can result in a system vulnerable to which type of attack?

Answers:

***Buffer overflow attack**

Denial-of-service attack

Dictionary attack

Superzapping

Explanation:

Poor software development practices and failing to program input validation checks can leave a system vulnerable to buffer overflow attacks. A buffer overflow occurs when software code receives more input than it was designed to handle because the programmer of that code failed to include input validation checks. When a buffer overflow occurs, the extra data is pushed into the execution stack and processed with the security context of the system itself. In other words, a buffer overflow attack often allows the attacker to perform any operation on a system.

Denial-of-service attacks exploit vulnerabilities in implementation and coding errors.

Dictionary attacks are waged against logon prompts or stolen copies of a security account's database.

Superzapping attacks are specific attacks that use a specialized utility named superzap to bypass the security of IBM mainframes to perform system alterations.

q_webattk_buffer_04_secp8

Which type of attack is the act of exploiting a software program's free acceptance of input in order to execute arbitrary code on a target?

Answers:

Data diddling

***Buffer overflow attack**

TOCTOU

Covert channel exploitation

Explanation:

The act of exploiting a software program's free acceptance of input in order to execute arbitrary code on a target is called a buffer overflow.

Data diddling is the change or corruption of data.

TOC/TOU is a logon session replay attack.

Covert channel exploitation is the use of timing or storage mechanisms to bypass security controls in order to leak information out of a secured environment.

q_webattk_buffer_05_secp8

A software development company pushes a critical update for its operating system, addressing security vulnerabilities.

The chief information security officer (CISO) schedules a meeting with the security team to discuss the specifics of one of these vulnerabilities exploited in recent cyberattacks.

Based on common operating system vulnerabilities, which of the following has insufficient or missing data validation mechanisms that lead to the system interpreting unintended command execution?

Answers:

***Buffer overflow**

Privilege escalation

Side-channel attack

Fingerprinting

Explanation:

Buffer overflow vulnerabilities occur when an application receives more data than it can handle, resulting in the potential for unintended command execution.

Privilege escalation refers to an attacker gaining unauthorized privileges on a system, usually through exploiting system misconfigurations or vulnerabilities that allow them to elevate their permissions.

Side-channel attacks are techniques that gather information indirectly from the physical implementation of a computer system rather than from software vulnerabilities or direct intrusion attempts.

Fingerprinting is the method used to determine the operating system running on a remote host, which does not involve exploiting a specific vulnerability within the operating system.

q_webattk_buffer_06_secp8

An organization experienced disruptions in its business operations due to repeated crashes of a critical application.

The IT department suspects a security issue and initiates a thorough investigation. The findings suggest the application failed due to excessive data processing beyond its expected limits.

What does the IT department confirm is the cybersecurity threat occurring in its business operations?

Answers:

***Buffer overflow attack**

Distributed denial-of-service (DDoS) attack

On-path attack

Social engineering attack

Explanation:

Buffer overflow attacks occur when an application receives more data than it can process, which can cause the application to crash or allow an attacker to execute arbitrary code. This is consistent with the symptoms described in the scenario.

A DDoS attack overwhelms a system's resources with traffic but does not involve overflowing an application's buffer with excessive data.

An on-path attack involves an attacker intercepting and possibly altering the communication between two parties and does not align with the application crashes described.

Social engineering attacks manipulate people into divulging confidential information. They do not involve direct attacks on an application's memory or cause application crashes.

q_webattk_buffer_07_secp8

An attacker passes data that deliberately floods a temporary memory space on one of the corporate machines.

What type of application vulnerability does this action describe?

Answers:

***Buffer overflow**

Race conditions

Malicious update

Memory injection

Explanation:

To exploit a buffer overflow vulnerability, the attacker passes data that deliberately floods a temporary memory space.

Application race condition vulnerabilities refer to software flaws associated with the timing or order of events within a software program, which can cause undesirable or unpredictable outcomes through manipulation.

A malicious update is an update that appears legitimate but contains harmful code, often used by cybercriminals to distribute malware or execute a cyber attack.

Memory injection vulnerabilities refer to a security flaw where an attacker can introduce or inject malicious code into a running application's process memory.

q_webattk_canonicalization_attack_secp8

You are a security analyst at a large e-commerce company. During a routine security audit, you discover that the company's web application is vulnerable to a canonicalization attack. The application uses file paths provided by users to access and display files.

Which of the following actions should you recommend to mitigate this vulnerability?

Answers:

Implement SSL/TLS encryption for all data in transit.

***Implement a whitelist of approved file paths and reject any requests that do not match.**

Switch to a different programming language that is less susceptible to canonicalization attacks.

Disable all client-side scripting.

Explanation:

Implementing a whitelist of approved file paths and rejecting any requests that do not match is a key defense against canonicalization attacks and is the correct answer. This ensures that only approved paths can be accessed, preventing an attacker from accessing unauthorized files.

While SSL/TLS encryption is important for protecting data in transit, it does not prevent canonicalization attacks. Canonicalization attacks manipulate the interpretation of file paths or URLs, which can occur regardless of whether the data is encrypted.

Simply switching to a different programming language does not inherently make the application more secure. If the application does not properly handle file paths or URLs, it could still be vulnerable to canonicalization attacks regardless of the programming language used.

Disabling all client-side scripting would not prevent canonicalization attacks. Canonicalization attacks target the server-side processing of file paths or URLs, not client-side scripts.

q_webattk_cross_secp8

Which of the following is an attack that injects malicious scripts into web pages to redirect users to fake websites to gather personal information?

Answers:

***XSS**

SQL injection

DLL injection

Drive-by download

Explanation:

Cross-site scripting (XSS) is an attack that injects scripts into web pages. When a user views the web page, the malicious scripts run, allowing the attacker to capture information or perform other actions.

XSS often relies on social engineering or phishing to entice users to click on links to web pages that contain malicious scripts.

Some scripts redirect users to legitimate websites but run in the background to capture information sent to the legitimate site.

Scripts can be written to read (steal) cookies that contain identity information (such as session information).

Scripts can also be designed to run under the security context of the current user. For example, scripts might execute with full privileges on the local system, or the scripts might run using the credentials used on a financial website.

A drive-by download is an attack where software or malware is downloaded and installed without explicit consent from the user.

An SQL injection attack occurs when an attacker includes database commands within user data input fields on a form, and those commands subsequently execute on the server.

A DLL injection attack occurs when a program is forced to load a dynamic-link library (DLL). This DLL then executes under the security context of the running application and executes malicious code included with the injected DLL.

q_webattk_csrf_secp8

You are a security analyst for a large e-commerce company. You notice that a significant number of users are complaining about unauthorized transactions on their accounts.

After investigating, you suspect that the site might be a victim of a Cross-Site Request Forgery (CSRF) attack.

Which of the following actions should you take next?

Answers:

Ignore the complaints as false positives since the transactions are coming from the users' IP addresses.

Implement CAPTCHA on all transaction pages to prevent automated attacks.

***Implement a CSRF token in your web application's forms.**

Advise users to change their passwords immediately.

Explanation:

Implementing a CSRF token in your web application's forms is a common and effective defense against CSRF attacks and is the correct answer. The token is a random value associated with the user's session and is included in every form submission. If the token in the form doesn't match the token on the server, the request is rejected.

Ignoring the complaints is not advisable as CSRF attacks can be performed from the users' browsers, making it appear as if the transactions are coming from their IP addresses.

While CAPTCHA can help prevent automated attacks, it is ineffective against CSRF attacks. CSRF attacks involve tricking the user into submitting a request, so a CAPTCHA would not prevent this.

While advising users to change their passwords is generally a good security measure, it would

q_webattk_injection_secp8

A cybersecurity analyst is trying to discover why a web application's interpreter executes unintended commands. The analyst notices that the attack involves sending untrusted data to an interpreter as part of a command or query.

What does the analyst suspect is happening?

Answers:

***Injection attack**

Replay attack

On-path attack

DDoS attack

Explanation:

Injection attacks involve sending untrusted data to an interpreter as part of a command or query. This data tricks the interpreter into executing unintended commands, potentially allowing unauthorized access or data retrieval.

Replay attacks involve the malicious repetition or delayed transmission of valid data. This type of attack does not relate to manipulating interpreters with untrusted data.

An on-path attack occurs when an attacker intercepts and possibly modifies packets in transit. It does not involve sending untrusted data to an interpreter as part of a command or query.

Distributed denial-of-service (DDoS) attacks overwhelm systems with traffic, aiming to make them unavailable. These attacks do not involve sending untrusted data to an interpreter.

q_webattk_input_secp8

Which of the following is specifically meant to ensure that a program operates on clean, correct, and useful data?

Answers:

***Input validation**

Application hardening

Process spawning

Error and exception handling

Explanation:

Input validation is the process of ensuring that a program operates on clean, correct, and useful data. Input validation uses routines (also called validation rules or check routines) that check for correctness, meaningfulness, and secureness in data input to the system.

Application hardening is the process of preventing vulnerability exploitation in software applications.

Error and exception handling is a programming language construct designed to handle the occurrence of exceptions (which are special conditions that change the normal flow of program execution).

Process spawning is the creation of a new process (also called a child process) by an existing process (also called a parent process).

q_webattk_ldap_injection_secp8

You are a security engineer at a large corporation. During a routine security audit, you discover that the company's internal application is vulnerable to Lightweight Directory Access Protocol (LDAP) Injection attacks.

The application uses LDAP for user authentication and does not sanitize input when adding filter parameters.

Which of the following actions should you recommend to mitigate this vulnerability?

Answers:

Implement SSL/TLS encryption for all data in transit.

Switch to a different authentication protocol like OAuth.

***Implement proper input validation and sanitization.**

Disable all client-side scripting.

Explanation:

Implementing proper input validation and sanitization is a key defense against LDAP Injection attacks and is the correct answer. This includes validating and sanitizing all user inputs and also using techniques like parameterized queries or prepared statements to ensure that user input is not interpreted as code.

While SSL/TLS encryption is important for protecting data in transit, it does not prevent LDAP Injection attacks. LDAP Injection attacks manipulate the structure of LDAP queries, which can occur regardless of whether the data is encrypted.

Simply switching to a different authentication protocol like OAuth does not inherently make the application more secure. If the application does not properly sanitize input, it could still be vulnerable to other types of injection attacks.

Disabling all client-side scripting would not prevent LDAP Injection attacks. LDAP Injection attacks target the server-side processing of LDAP queries, not client-side scripts.

q_webattk_memory_injection_secp8

You are a software engineer developing a new application for a healthcare provider.

During a code review, you notice that the application is using unmanaged code and directly accessing memory addresses, which could potentially lead to memory injection attacks.

Which of the following actions should you take next?

Answers:

Continue with the current design, but implement strong input validation.

***Switch to a managed code language that handles memory management automatically.**

Implement a custom memory management system to prevent unauthorized access.

Use encryption for all data stored in memory.

Explanation:

Switching to a managed code language that handles memory management automatically can significantly reduce the risk of memory injection attacks and is the correct answer. Managed code languages like Java or C# provide built-in protections against direct memory access.

While strong input validation is a good security practice, it alone may not prevent memory injection attacks. Attackers can often craft malicious payloads that bypass common input validation techniques.

Implementing a custom memory management system can be complex and error-prone. It's generally better to rely on the memory management provided by modern programming languages, which have been extensively tested and refined.

While encryption can protect data at rest and in transit, it does not prevent memory injection attacks, which exploit the way an application accesses and manipulates memory.

q_webattk_privilege_escalation_secp8

You are a security engineer at a large corporation. During a routine security audit, you discover that the company's internal application allows users to view the profiles of other users by manipulating the URL. This could potentially lead to a privilege escalation attack.

Which type of privilege escalation is this, and what should you do to mitigate this vulnerability?

Answers:

This is a vertical privilege escalation. Implement role-based access control to prevent users from accessing higher-level privileges.

***This is a horizontal privilege escalation. Implement proper input validation and sanitization.**

This is a vertical privilege escalation. Implement proper input validation and sanitization.

This is a horizontal privilege escalation. Implement role-based access control to prevent users from accessing other users' profiles.

Explanation:

This is a horizontal privilege escalation, and implementing the proper input validation and sanitization is the correct answer. Horizontal privilege escalation involves accessing other users' profiles or data. Implementing proper input validation and sanitization can help prevent users from manipulating the URL to access other users' profiles.

Vertical privilege escalation involves elevating a user's privileges to a higher level, not accessing other users' profiles. Implementing role-based access control is a good security measure, but it may not prevent horizontal privilege escalation.

Vertical privilege escalation involves elevating a user's privileges to a higher level, not accessing other users' profiles. Implementing proper input validation and sanitization is a good security measure, but it does not address the type of privilege escalation in question.

While implementing role-based access control is a good security measure, it may not prevent horizontal privilege escalation if users can still manipulate the URL to access other users' profiles. Proper input validation and sanitization is needed to address this vulnerability.

q_webattk_race_conditions_secp8

Which type of application vulnerability can refer to software flaws associated with the timing or order of events within a software program, which can cause undesirable or unpredictable outcomes through manipulation?

Answers:

***Race conditions**

Memory injection

Buffer overflow

Malicious update

Explanation:

Application race condition vulnerabilities refer to software flaws associated with the timing or order of events within a software program, which can cause undesirable or unpredictable outcomes through manipulation.

Memory injection vulnerabilities refer to a type of security flaw where an attacker can introduce (inject) malicious code into a running application's process memory.

To exploit a buffer overflow vulnerability, the attacker passes data that deliberately floods a space of temporary memory.

A malicious update refers to an update that appears legitimate but contains harmful code, often used by cybercriminals to distribute malware or execute a cyberattack.

q_webattk_replay_attack_secp8

You are a network security engineer at a large corporation. During a routine security audit, you discover that the company's internal network is vulnerable to replay attacks. The network uses a simple challenge-response authentication protocol.

Which of the following actions should you recommend to mitigate this vulnerability?

Answers:

Implement SSL/TLS encryption for all data in transit.

***Implement a time-based one-time password (TOTP) system.**

Switch to a different authentication protocol like OAuth.

Disable all client-side scripting.

Explanation:

Implementing a time-based one-time password (TOTP) system can help prevent replay attacks and is the correct answer. In a TOTP system, the authentication token changes every few seconds, making captured authentication messages useless after a very short period of time.

While SSL/TLS encryption is important for protecting data in transit, it does not prevent replay attacks. Replay attacks involve capturing and retransmitting authentication messages, which can occur regardless of whether the data is encrypted.

Simply switching to a different authentication protocol like OAuth does not inherently make the network more secure. If the new protocol does not include protections against replay attacks, the network could still be vulnerable.

Disabling all client-side scripting would not prevent replay attacks. Replay attacks target the authentication protocol, not client-side scripts.

q_webattk_sql_01_secp8

Which of the following is subject to SQL injection attacks?

Answers:

***Database servers**

Web servers serving static content

Browsers that allow client-side scripts

ActiveX controls

Explanation:

A SQL injection attack occurs when an attacker includes database commands within user data input fields on a form, which subsequently execute on the server. The injection attack succeeds if the server does not properly validate the input to restrict the entry of characters that could end and begin a database command.

The following do not directly rely on database servers and are not subject to SQL injection attacks:

Web servers serving static content

Browsers that allow client-side scripts

ActiveX controls

q_webattk_sql_02_secp8

You have a website that accepts input from users for creating customer accounts. Input on the form is passed to a database server where the user account information is stored.

An attacker is able to insert database commands in the input fields and have those commands execute on the server.

Which type of attack has occurred?

Answers:

***SQL injection**

Buffer overflow

DLL injection

Cross-site scripting

Explanation:

A SQL injection attack occurs when an attacker includes database commands within user data input fields on a form, which subsequently execute on the server. The injection attack succeeds if the server does not properly validate the input to restrict the entry of characters that could end and begin a database command. SQL injection attacks are prevented by proper programming methods that prevent commands from occurring within form data or that filter data to prevent such attacks.

A buffer overflow occurs when an operating system or application does not properly enforce boundaries for how much and which type of data can be inputted. Hackers submit data beyond the size reserved for the data in the memory buffer, and the extra data overwrites adjacent memory locations. The extra data sent by the attacker could include executable code that might then be able to execute in privileged mode.

Cross-site scripting (XSS) is an attack that injects scripts into web pages. When the user views the web page, the malicious scripts run, allowing the attacker to capture information or perform other actions.

A DLL injection attack occurs when a program is forced to load a dynamic-link library (DLL). This DLL then executes under the security context of the running application and executes malicious code included with the injected DLL.

q_webattk_sql_03_secp8

An attacker inserts SQL database commands into a data input field of an order form used by a web-based application. When submitted, these commands are executed on the remote database server, causing customer contact information from the database to be sent to the malicious user's web browser.

Which practice would have prevented this exploit?

Answers:

***Implementing client-side validation**

Using the latest browser version and patch level

Installing antivirus, anti-spyware, pop-up blockers, and firewall software

Implementing a script blocker

Explanation:

Client-side validation should have been used on the local system to identify input errors in the order form before the data was sent to the server. In this example, if the user entered SQL commands in an order form field, the error would have been immediately detected and blocked before the data was submitted to the server.

Using the latest browser version and patch level, installing anti-malware software, and using a script blocker are valuable security measures. But these would not have prevented the exploit in this scenario.

q_webattk_sql_04_secp8

In 2011, Sony was targeted by an SQL injection attack that compromised over one million emails, usernames, and passwords.

Which of the following could have prevented the attack?

Answers:

***Careful configuration and penetration testing on the front end**

Scanning the operating system and application regularly for bugs and errors

Using VPN technology to protect client data when connecting from a remote system

Blocking, or at least monitoring, activity on ports 161 and 162

Explanation:

SQL attacks such as with Sony, United States Department of Energy, and MySQL could have been prevented with careful configuration and penetration testing on the front end.

One of the steps for preventing privilege escalation is to scan the operating system and application regularly for bugs and errors.

To defend against WPA/WPA2 cracking, use VPN technology to protect client data when connecting from a remote system.

The easiest way to prevent SNMP exploitation is to block, or at least monitor, activity on ports 161, 162, and any other port you've configured for SNMP traffic.

q_webattk_sql_05_secp8

SQL injections are a result of which of the following flaws?

Answers:

Web applications

***The database**

The web server

The file system

Explanation:

SQL injections happen when a SQL statement is entered in a data field and gets executed in the database.

Web applications, web servers, and the file system are not databases and are not directly impacted by SQL injections.

q_webattk_sql_06_secp8

Which of the following functions does a single quote (') perform in an SQL injection?

Answers:

***Indicates that data has ended and a command is beginning.**

Indicates that code is ending and a comment is being entered.

Indicates that everything after the single quote is a comment.

Indicates that the comment has ended and data is being entered.

Explanation:

A single quote (') indicates that data has ended and a command is beginning.

The double dashes (--) indicate that code is ending and a comment is being entered.

Comments are code that a program does not execute and are usually used for explanations or reminders for the coder. Applications know to ignore the comments.

q_webattk_ssrf_secp8

You are a security engineer at a large financial institution. You notice some unusual network traffic patterns indicating that internal resources are being accessed from the public-facing web server.

You suspect a Server-Side Request Forgery (SSRF) attack.

Which of the following actions should you take next?

Answers:

Implement rate limiting on the public-facing web server.

Implement input validation and sanitization on all user inputs.

Block all incoming traffic to the public-facing web server.

***Restrict the public-facing web server from making unnecessary outbound requests.**

Explanation:

One of the most effective ways to mitigate SSRF attacks is to restrict the public-facing web server from making unnecessary outbound requests. This can be achieved by implementing network-level controls to block outbound traffic from the server to internal resources or by configuring the server or application to deny all outbound requests unless they are absolutely necessary.

While rate limiting can help mitigate some types of attacks (like DDoS), it would not be effective against an SSRF attack. SSRF attacks do not rely on a high volume of requests but rather on exploiting the trust relationship between the server and other internal systems.

While input validation and sanitization are good security practices, they alone may not prevent SSRF attacks. Attackers can often craft malicious payloads that bypass common input validation techniques.

Blocking all incoming traffic to the public-facing web server would effectively take your website offline, which is not a practical solution.

q_webattk_web_secp8

While using a web-based order form, an attacker enters an unusually large value in the Quantity field.

The value they entered is so large that it exceeds the maximum value supported by the variable type used to store the quantity in the web application. This causes the value of the quantity variable to wrap around to the minimum possible value, which is a negative number.

As a result, the web application processes the order as a return instead of a purchase, and the attacker's account is credited with a large sum of money.

Which practices would have prevented this exploit? (Select two.)

Answers:

***Implementing client-side validation**

Using the latest browser version and patch level

Installing antivirus, anti-spyware, pop-up blockers, and firewall software

Installing the latest operating system updates

***Implementing server-side validation**

Explanation:

Client-side validation and server-side validation should have been used to identify input errors in the order form. In this example, if the user entered an invalid quantity in an order form field, client-side validation would have detected and blocked the error before the data was submitted to the server. Server-side validation should have also been used after the data was sent to the server to detect errors. Experienced attackers can circumvent client-side validation techniques by sending data to the server from outside the application's standard user interface, bypassing any input validation measures that may have been implemented on the client.

Using the latest browser version and patch level, installing the latest operating system updates, and using a script blocker are valuable security measures, but they would not have prevented the exploit in this scenario.

q_webattk_xml_injection_secp8

You are a security analyst at a software company. During a security audit, you discover that the company's web application is vulnerable to Extensible Markup Language (XML) Injection attacks.

The application uses XML to transport data between the client and the server.

Which of the following actions should you recommend to mitigate this vulnerability?

Answers:

Implement SSL/TLS encryption for all data in transit.

Use a less complex data format like JSON instead of XML.

***Implement proper input validation and sanitization.**

Disable all client-side scripting.

Explanation:

Implementing proper input validation and sanitization is a key defense against XML Injection attacks and is the correct answer. This includes validating and sanitizing all user inputs and also using techniques like parameterized queries or prepared statements to ensure that user input is not interpreted as code.

While SSL/TLS encryption is important for protecting data in transit, it does not prevent XML Injection attacks. XML Injection attacks manipulate the structure of XML data, which can occur regardless of whether the data is encrypted.

Simply switching to a less complex data format like JSON does not inherently make the application more secure. JSON is also susceptible to injection attacks (JSON Injection) if proper security measures are not in place.

Disabling all client-side scripting would not prevent XML Injection attacks. XML Injection attacks target the server-side processing of XML data, not client-side scripts.

sql_attack_question1

What is Blake Jackson's account balance?

Answers:

$582.29

582.29

sql_attack_question2

What is the account number of Nisha Dickson?

Answers:

90003

8.9 Application Development and Security

As you study this section, answer the following questions:

- What are two common standardized software development models?
- How should security be implemented in the different stages of development?
- What are the responsibilities of developers after a product is released?
- What are some important application hardening techniques?

In this section, you will learn to:

- Harden applications on Linux.
- Implement application whitelisting with AppLocker.
- Implement Data Execution Preventions (DEP).

The key terms for this section include:

Term	Definition

Normalization	Data reorganized in a relational database to eliminate redundancy by having all data stored in one place and storing all related items together.
Stored procedures	One or more database statements stored as a group in a database's data dictionary, which when called, executes all the statements in the collection.
Code obfuscation	The deliberate act of creating source or machine code that is difficult for humans to understand. In other words, the code is camouflaged.
Code reuse	Using the same code multiple times.
Dead code	Code that is non-executable at run-time, or source code in a program that is executed but is not used in any other computation.
Memory management	A resource management process applied to computer memory. It allows your computer system to assign portions of memory, called blocks, to various running programs to optimize overall system performance.
Third-party libraries	A library where the code is not maintained in-house.
Software Development Kits (SDKs)	A set of software development tools that can be installed as one unit.
Data exposure	Unintended exposure of personal and confidential data.
Fuzz testing	A software testing technique that exposes security problems by providing invalid, unexpected, or random data to the inputs of an application.
Code signing	The process of digitally signing (encrypting) executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.3 Explain the importance of change management processes and the impact to security. Version control

2.2 Explain common threat vectors and attack surfaces.

- Unsupported systems and applications

2.5 Explain the purpose of mitigation techniques used to secure the enterprise.

- Application allow list
- Patching
- Monitoring
- Hardening techniques
  - Installation of endpoint protection
  - Host-based firewall
  - Removal of unnecessary software

3.3 Compare and contrast concepts and strategies to protect data.

- Methods to secure data
  - Obfuscation

4.1 Given a scenario, apply common security techniques to computing resources.

- Secure baselines
  - Establish
- Application security
  - Input validation
  - Secure cookies
  - Static code analysis
  - Code signing
- Sandboxing

4.3 Explain various activities associated with vulnerability management.

- Identification methods
  - Application security

Static analysis

Dynamic analysis

4.5 Given a scenario, modify enterprise capabilities to enhance security.

Operating system security

SELinux

4.6 Given a scenario, implement and maintain identity and access management.

Provisioning/de-provisioning user accounts

Access controls

Least privilege

4.7 Explain the importance of automation and orchestration related to secure operations.

Use cases of automation and scripting

User provisioning

Resource provisioning

Security groups

Ticket creation

Enabling/disabling services and access

Integrations and Application programming interfaces (APIs)

Benefits

Efficiency/time saving

Enforcing baselines

Staff retention

Other considerations

Complexity

Cost

	Single point of failure Technical debt Ongoing supportability 5.1 Summarize elements of effective security governance. Policies Software development lifecycle (SDLC) 5.2 Explain elements of the risk management process. Risk analysis Exposure factor
TestOut Security Pro	3.2 Implement application defenses 3.2.1 Implement an application allow list

8.9.1 Development Life Cycle (Lesson Video)

Transcript:

In this lesson, we'll look at how to apply security concepts to software development. Even though you may not be a software developer yourself, most security professionals work with software engineers in some capacity. Just like other company resources, the applications they develop need to be properly secured.

Today, we'll look at the two most common development lifecycles to understand how you can secure each step throughout the development process.
The most widely used development model is the Waterfall model. It's called this because each step is completed before the next step is begun. Each step flows to the next like a waterfall.
The first step in a Waterfall model is Requirements. All requirements for the application need to be gathered from the client, user, or stakeholder. The next step is Design. This is when the software is documented, diagrammed, and designed. After that, we have Implementation. This is the actual coding and building of the application.
We then have the Testing phase. During this phase, a quality assurance team makes sure that the team met the requirements, that the code works properly across devices, and that security flaws and vulnerabilities were phased out.
After Testing comes the Deployment. This is when the application is released to a client or to the public. The final step is Maintenance. The application is in use in this phase, but it's monitored for bugs or problems that are quickly patched. This is an ongoing stage that continues throughout the app's life.
Understand that the application will likely go through some of these steps multiple times before moving on. For example, the application might go through the Design step five times before it's ready to move on to the Implementation step. Then the application may move from the Implementation stage back to the Design stage if a new feature needs to be added.
This entire development cycle is a slow process and may take months or years to complete. The Waterfall method lacks flexibility since the requirements determined in the beginning carry through to the end product.
A more agile approach was introduced in 2001 which approaches software development as a continuous, changing process with never-ending versions, bug fixes, and enhancements. This approach is aptly named Agile. One reason the Agile methodology is pragmatic is because security vulnerabilities are constantly addressed with new updates and fixes.

Unlike Waterfall, Agile doesn't work on the entire application at once. Instead, it breaks development into smaller time frames called Sprints. Each Sprint has a specific duration, usually two to three weeks. Within this time frame, developers work on a specific application feature.
These features typically go through the same stages that all applications go through, such as Requirements, Design, Development, Testing, and Deployment. At the end of the Sprint, the developers move on to the next feature.
With Agile, you perform testing throughout the development cycle to help catch security vulnerabilities early. Waterfall, on the other hand, leaves testing until the end of the cycle. This can cause testing time to be cut short to meet deadlines.
Coding errors and design flaws are the main causes of software vulnerabilities. We can categorize these errors into two types: compile errors and runtime errors.
A compile error refers to an error that occurs during the building or compilation stage. The error compromises the software implementation, which prevents it from running at all.
A runtime refers to an error that occurs while the software is running. The software begins execution, but it fails when it cannot resolve a problem it encounters.
During the coding and design phases, you can increase application security by integrating security testing into each step of the process. Let's look at three testing methods: static, dynamic, and interactive application security testing.
Static application security testing, or SAST, is also known as white box testing. SAST focuses on analyzing source code, binaries, and byte code early in the development process. SAST tools are good at identifying things like SQL-injection vulnerabilities and buffer overflows. They can identify the exact cause of a coding problem, but only in code not yet deployed. They're language specific, but you can run them continually and apply them widely. Please note that they do have a high percent of false positives and are limited in the types of vulnerabilities they can detect.
Dynamic application security testing, or DAST, is also known as black box testing. DAST scans a deployed application once it enters runtime. The results are based on how the application responds to a series of tests from the outside. These tests aren't language specific and have a lower rate of false positives. The downsides to DASTs are that they're hard to automate, can't pinpoint the exact cause of a flaw, and can take up to a week to complete.
Finally, we have interactive application security testing, or IAST, which can be broken down into two types: passive and active. In passive IAST, we build interactivity into static application security testing. IAST tools are source code scanners that work during runtime.
In active IAST, the testing tools can access interpreters and compliers, allowing precise identification of a problematic line of code during runtime. This speeds up the testing and remediation process. The combination of passive and active IAST can help us in the development stage by catching vulnerabilities early, in the QA stage by adding automated security checkpoints, and in the production stage through continuous monitoring.
That's it for this lesson. In this lesson, we reviewed two common software development lifecycles: Waterfall and Agile. Then we looked at two different kinds of coding errors: compile errors and runtime errors. We finished up by looking at static, dynamic, and interactive application security testing and how these tools can help us secure applications throughout the development lifecycle.

8.9.2 Automation and Scripting (Lesson Video)

Transcript:

Automation is a powerful tool for managing security operations. Automation allows us to perform repetitive, rule-based tasks to improve efficiency and reduce user errors. Combining automation systems with orchestration and scripting can make a huge difference in the efficiency of security operations. In this lesson, we'll go over these systems and then look at the benefits and considerations of automation, orchestration, and scripting.

Automation uses software to perform repetitive, rule-based tasks, such as monitoring for threats, applying patches, maintaining baselines, or responding to incidents, to improve efficiency and reduce the likelihood of human error.
We can develop scripts to automate many of the common repetitive tasks. Scripts tell the system what steps to take to automate the task. Orchestration enhances automation by coordinating and streamlining the interactions between automated processes and systems. Orchestration supports seamless and integrated workflows, especially in large, complex environments with many different security tools and systems.
Scripting uses a specific language, such as Python, JavaScript, or PowerShell, to perform tasks on a computer or network system. A script is typically a text file with commands written that tell the system what to do step by step.
Typical programs must be compiled before running, which means the code's text must be converted into binary code. Scripts, on the other hand, are read and executed line by line. Scripting languages are usually something a human can

look at and interpret, whereas binary code is much more difficult to read. When combined, scripting and automation are critical tools in modern IT operations.

For all of these different platforms and tools to communicate with each other, we need to implement the appropriate Application Programming Interface, or API. An API simply serves as an interpreter between multiple systems so they can all speak the same language. Automation and orchestration systems will need an appropriate API so they can be integrated into the network systems.

Automation can help an organization to streamline processes, enhance security, and improve efficiency. For example, we can develop scripts to apply security policies on network devices and use automation software to help enforce the security policies. Another example is using scripts to apply patches and updates across the organization and automation tools to track the changes and notify us if anything goes wrong.

Using automation not only enhances the efficiency of security operations but also reduces the burden on the security team and reduces the likelihood of human error. Security analysts must monitor numerous systems for potential threats, manage high volumes of alerts, including many false positives, and respond to confirmed threats as quickly as possible. These working conditions often lead to long hours, anxiety, and elevated stress levels, resulting in operator fatigue. This fatigue is a significant concern in cybersecurity because it can lead to decreased alertness and cognitive function and impair the ability of security personnel to identify and respond to threats effectively. Fatigue results in missed critical alerts, slower response times, and a greater likelihood of errors, any of which can compromise security.

Having to perform common repetitive tasks, such as user and resource provisioning, assigning security groups, creating tickets, and other common tasks, can be minimized by creating scripts and monitoring with automation tools. This will greatly help reduce operator fatigue, leading to fewer errors and better security operations.

While implementing automation can make the job easier and more efficient, there are some challenges we need to be aware of. Appropriate automation and orchestration solutions can be extremely complex and require a deep knowledge and understanding of the organization's systems, processes, and interdependencies. A poorly planned or executed strategy will make systems more difficult to manage and maintain.

When implementing automation solutions, the initial cost can be quite high. Even though the organization will save money in the long term, the high initial cost can make it difficult to implement. If a critical automated system or process fails, this can impact multiple areas of the organization. Having this single point of failure is something that should be considered, and proper mitigation steps should be implemented.

If automation and orchestration systems are poorly designed and implemented, this can lead to system instability, complexity, and increased cost. This is known as technical debt. Lastly, automation systems require ongoing support and updates to stay effective and secure. Without proper support, these systems will be quickly outdated.

That'll wrap up this lesson on automation and scripting. In this lesson, we first went over what automation, orchestration, and scripting are. We then looked at the benefits of using these technologies, including reducing operator fatigue and increasing the efficiency of our security operations. We wrapped up by going over some concerns that should be considered when implementing automation orchestration solutions.

8.9.3 SDLC and Development Facts

Even though you may not be a software developer, most security professionals at some point work with software engineers who develop applications. And just like any other enterprise component, these applications need to be properly secured.

This lesson covers the following topics:

- Waterfall development life cycle model
- Agile development life cycle model
- Coding errors
- Error handling
- Static code analysis
- Software sandboxing

Waterfall Development Life Cycle Model

The most widely used development model is the Waterfall model. It is called this because each step is completed before the next step is begun so that each step flows to the next.

The Waterfall development life cycle model steps are:

Step	Description
Requirements	All requirements for the application being developed are gathered from the client, user, or stakeholder.
Design	The software is documented, diagramed, and designed.
Implementation	The code is written.
Testing	A quality assurance team makes sure requirements are met, the code works properly across devices, and security issues are noted.
Deployment	The application is released to a client or the public.
Maintenance	The application is monitored for bugs or problems that are patched or fixed while in use. This is an ongoing stage that continues throughout the life of the app.

Understand that an application will likely go through some of these steps multiple times before moving to the next step. For example, the application might go through the Design step five times before it's ready to move to the Implementation step. Or the application may move back from the Implementation stage to the Design stage if a new feature needs to be added.

This entire development life cycle is a slow process and may take months or years to complete. The Waterfall method also lacks flexibility since the requirements determined in the beginning are carried through to the end product.

Agile Development Life Cycle Model

A more agile approach was introduced in 2001 that approaches software development as a continuous, changing process with never-ending versions, bug fixes, and enhancements. This approach is aptly named Agile.

The Agile model works in this manner:

- It breaks development into smaller time frames called Sprints.
  - Each Sprint has a specific duration (usually two to three weeks).
  - Developers work on one feature during a Sprint.
- At the end of each Sprint, developers move on to the next feature.
- Testing is performed throughout the development cycle.

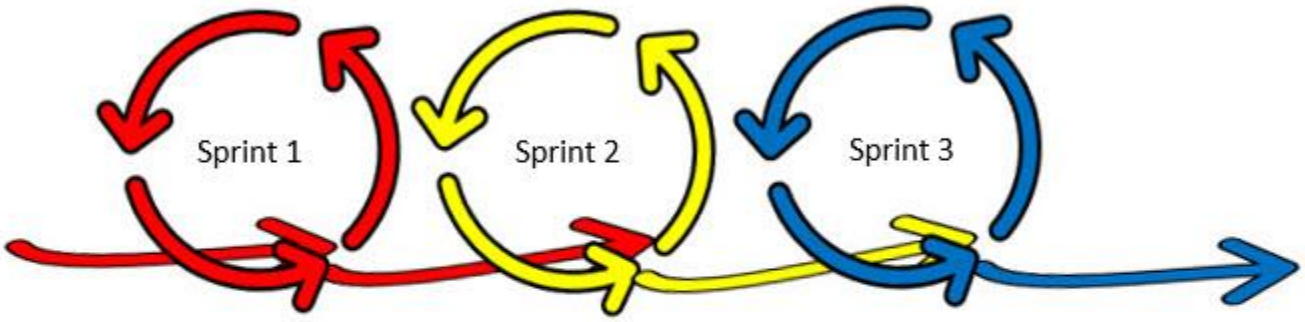

1. Requirements	1. Requirements	1. Requirements
2. Design	2. Design	2. Design
3. Development	3. Development	3. Development
4. Testing	4. Testing	4. Testing
5. Deployment	5. Deployment	5. Deployment

Coding Errors

Coding errors and design flaws are the main causes of software vulnerabilities. We can categorize these errors into two types:

Coding Error Type	Description
Compile	An error that occurs during the building or compilation stage Error compromises the software implementation Prevents the app from running
Runtime	An error that occurs while software is running Sometimes called bugs

Error Handling

A well-written application must be able to handle errors and exceptions gracefully. This means that the application performs in a controlled way when something unpredictable happens. An error or exception could be caused by invalid user input, a loss of network connectivity, another server or process failing, etc. Ideally, the programmer will have written a structured exception handler (SEH) to dictate what the application should do. Each procedure can have multiple exception handlers.

Some handlers will deal with anticipated errors and exceptions; there should also be a catchall handler that will deal with the unexpected. The main goal must be for the application not to fail in a way that allows the attacker to execute code or perform some sort of injection attack. One infamous example of a poorly written exception handler is the Apple GoTo bug (nakedsecurity.sophos.com/2014/02/24/anatomy-of-a-goto-fail-apples-ssl-bug-explained-plus-an-unofficial-patch).

Another issue is that an application's interpreter may default to a standard handler and display default error messages when something goes wrong. These may reveal platform information and the inner workings of code to an attacker. It is better for an application to use custom error handlers so that the developer can choose the amount of information shown when an error is caused.

Technically, an error is a condition that the process cannot recover from, such as the system running out of memory. An exception is a type of error that can be handled by a block of code without the process crashing. Note that exceptions are still described as generating error codes/messages, however.

Static Code Analysis

Static code analysis is a crucial software development practice. It involves scrutinizing source code to identify potential vulnerabilities, errors, and non-compliant coding practices before the program is finalized. By examining code in a 'static' state, developers can catch and rectify issues early in the development lifecycle. This makes it a proactive approach to building secure, reliable, high-quality software.

Application security approaches focus on software development and deployment lifecycles, with a heavy emphasis on secure coding practices that encourage developers to write code that prevents common vulnerabilities like SQL injection and cross-site scripting. Application security practices also mandate static and dynamic application security testing (SAST). Coding practices designed to support regular patching and updates are crucial to support the prompt resolution of newly discovered vulnerabilities.

Static code analysis supports secure coding and is performed using specialized tools, often integrated into software development suites. These tools automate code checks against pre-determined rules and flag potential issues so developers can review and address them. Some commonly used static analysis tools include SonarQube (https://www.sonarsource.com/products/sonarqube/), Coverity (https://www.synopsys.com/software-integrity/security-testing/static-analysis-sast.html), and Fortify (https://www.microfocus.com/en-us/cyberres/application-security), but there are many others.

Static code analysis in software development is critical because it enables early detection of bugs and security vulnerabilities and helps prevent potentially catastrophic failures in the final product. It also improves code quality and maintainability by enforcing coding standards and best practices. Additionally, static code analysis helps educate developers about common coding errors and security risks, which helps promote security-conscious development practices.

During the coding and design phases of development, you can increase development and application security by implementing a few key practices. Integrate security testing into each step in the development process.

Secure Testing Method	Description
Static application security testing	Known as white box testing Focuses on analyzing source code, binaries, and byte code early in the development process Good at identifying things like SQL injections and buffer overflows Can identify the exact cause of a coding problem: Only in code that's written but not deployed

	Is language specific Can run continually and be widely applied Has a high percentage of false positives Limited in the types of vulnerabilities it can detect
Dynamic application security testing	Known as black box testing Scans applications after deployment Tests from the outside Uses a series of tests to determine vulnerabilities and flaws Not language specific Has fewer false positives Hard to automate Cannot pinpoint the cause of a flaw Can take up to a week to complete the testing process
Interactive application security testing	Has two types: Passive: Interactive functionality is built into static application security testing. Uses source code scanners during runtime. Active: Testing tools can access interpreters and compilers, allowing precise identification of a problematic line of code in runtime. Speeds up testing and remediation. Can help in the Development stage by catching vulnerabilities early. Can help in the QA stage by adding automated security checkpoints.

	Can help in the Production stage through continuous monitoring.

Software Sandboxing

Sandboxing is a security mechanism used in software development and operation to isolate running processes from each other or prevent them from accessing the system they are running on. A sandbox is a protection feature designed to control a program so it runs with highly restrictive access. This containment strategy reduces the potential impact of malicious or malfunctioning software, making it effective for improving system security and stability and mitigating risks associated with software.

A practical example of sandboxing is implemented in modern web browsers, like Google Chrome, which separates each tab and extension into distinct processes. If a website or browser extension in one browser tab attempts to run malicious code, it is confined within that tab's sandbox. This action prevents malicious code from impacting the entire browser or underlying operating system. Similarly, if a tab crashes, it doesn't cause the whole browser to fail, improving reliability.

Operating systems also utilize sandboxing to isolate applications. For example, iOS and Android use sandboxing to limit each application's actions. An app in a sandbox can access its own data and resources but cannot access other app data or any nonessential system resources without explicit permission. This approach limits the damage caused by poorly written or malicious apps.

Virtual machines (VMs) and containers like Docker offer another example of sandboxing at a larger scale. Each VM or container can run in isolation, separated from the host and each other. The others remain unaffected if one VM or container experiences a security breach or system failure.

8.9.4 Version Control Management (Lesson Video)

Transcript:

When you develop an application, the most valuable possession is the source code. Having it protected is critical. In this lesson, we'll discuss version control management as well as provisioning and deprovisioning, which aid in protecting the source code. In application development, there's usually a complex process of writing, revising, and updating the code since multiple developers work together on the same project. There are different ways to approach this reality as this creates a challenge when keeping track of changes in the code. In general, most organizations use a version control system, or VCS, to track these changes.

Let's look at how it works. A version control system uses a repository, which is a storage location that holds all the source files used during development. This centralized file system can be accessed by all authorized users. It allows developers to simultaneously work on the same file, revert to older files, and even restore files that were deleted. These repositories are often stored in the cloud or on a third-party website.
There are many benefits to using a VCS. For one, it has the flexibility to allow for branching and merging. This allows a developer to work on one specific feature in a branch that he or she created, and then merge it back into the source code when ready. It also keeps a history of code changes. This history includes the details of who, what, when, and why. Together with annotations that developers can add, a team can understand why the code was designed the way it was. This also helps prevent incompatibility issues without preventing the developers from continuing their work.
Now let's talk about provisioning. Provisioning is the process of giving access to users through privileges and permissions. It's important that developers and others working on a project have access to all the resources they need to do their work. This includes the permissions that you give in the VCS as well as in other areas of your system. It's important to only give permissions for specifically what a user needs and only until he or she no longer needs that access. If he or she moves on to another project, changes roles, leaves the organization, or if the project is complete, be sure to deprovision them. This means that you remove their privileges and permissions. This is the basic idea behind the principle of least privilege, which is important in application development just as it is throughout your organization.

That's it for this lesson. In this lesson, we discussed important secure coding techniques that include using a version control system to provide continuity, traceability, and protection throughout the development process. We then went over provisioning and deprovisioning as they apply to protecting the source code.

8.9.5 Application Development Security Facts

In our world today, information is exchanged constantly. This means that attackers are working relentlessly to access our data. It is essential that we begin security efforts at the coding level.

This lesson covers the following topics:

- Secure coding techniques
- Code signing
- Secure cookies

Secure Coding Techniques

The security considerations for new programming technologies should be understood and tested before deployment. One of the challenges of application development is that the pressure to release a solution often trumps any requirement to ensure that the application is secure. A legacy software design process might be heavily focused on highly visible elements, such as functionality, performance, and cost. Modern development practices use a security development lifecycle running in parallel or integrated with the focus on software functionality and usability. Examples include Microsoft's SDL (microsoft.com/en-us/securityengineering/sdl) and the OWASP Software Assurance Maturity Model (owasp.org/www-project-samm). OWASP also collates descriptions of specific vulnerabilities, exploits, and mitigation techniques, such as the OWASP Top 10 (owasp.org/www-project-top-ten).

Secure coding concepts include the following:

Concept	Description
Normalization	Normalization is data reorganized in a relational database with the intent to eliminate redundancy by having all related data stored in one place. Normalization: Increases performance by reducing disk space. Provides quick and efficient access to manipulate the data. Lowers the risk of exploitation.
Stored procedures	Stored procedures are one or more database statements stored as a group in a database's data dictionary. When called, these procedures execute all the statements in the collection. Stored procedures: Centralize the code and eliminate the need to reproduce it. Keep calling program rules consistent across programs.

	Protect the code from users by allowing the user to call a stored procedure without seeing the actual code. Limit injection attacks.
Code obfuscation/code camouflage	Obfuscation is the deliberate act of creating source or machine code that is difficult for humans to understand. In other words, the code is camouflaged. Programmers use roundabout expressions to compose statements that deliberately obfuscate code to conceal its purpose or logic. They use implicit values embedded in it to prevent tampering, deter reverse engineering, or as a puzzle or recreational challenge for someone reading the source code. This is usually done with an automated tool but can also be done manually. There are several methods, but most can be reverse-engineered.
Code reuse	Code reuse is simply using the same code multiple times. Reusing code is a good idea if the programmer writes the same code at least three times. Code reuse: Can create a shared library for others that use the same code. Can be a security problem if: The code is not secure before it is shared and used multiple times. It has changes made to fit a new use, but the changes are not secure. Be sure to comprehensively test code before allowing reuse.
Dead code	Sometimes dead code refers to code that is non-executable at runtime. Sometimes, it means source code in a program that is executed but is not used in any other computation, making it obsolete. Remove any dead code from your application for security. If it does not exist, it cannot be exploited.
Memory management	Memory management is a resource-management process applied to computer memory. It allows your computer system to assign portions of memory called blocks to various running programs that optimize overall system performance. Many arbitrary code attacks depend on the target application having faulty memory management procedures. This allows the attacker to execute their own code in the space marked out by the target application. There are known unsecured practices for memory management that should be avoided, and checks for processing untrusted input, such as strings, to ensure that it cannot overwrite areas of memory.

	Memory management resides in the hardware, the operating system, programs, and applications. In the hardware, memory management involves components that physically store data, such as RAM chips, memory caches, and SSDs. In the OS, memory management involves the allocation of specific memory blocks to individual programs as user demands change. At the Application level, memory management ensures the availability of adequate memory for the objects and data structures of each running program at all times. When the program requests a block of memory, the allocator in the memory manager assigns that block to the program. When a program no longer needs the data in the previously allocated memory blocks, those blocks become available for reassignment. This task can be done automatically by the memory manager or manually by the programmer. The most common memory vulnerabilities include: Size of input in buffer copy not checked Buffer size was calculated incorrectly Format string not controlled To prevent vulnerabilities: Limit the amount of characters read into the buffer. Define constants for the size argument. Do not allow user input in format strings.
Third-party libraries and software development kits (SDKs)	A third-party library is a library where the code is not maintained in-house. A software development kit (SDK) is a set of software development tools that can be installed as one unit. Both can provide code frameworks or code snippets to help development go faster. Though they can be very helpful, there are risks involved. For example: Anytime code comes from an outside source, there is a risk that it may contain flaws and vulnerabilities Sometimes code comes in bundles, giving developers more code than they need Extra code can create extra opportunities for exploitation SDKs are often open-source and, as such, there may be no urgency to fix bugs Be sure to test code from third-party libraries and SDKs for functionality and security issues.
Sensitive data exposure	Sensitive data exposure involves unintended exposure of personal and confidential data. This can come from:

	Weak or missing encryption Coding flaws Misapplied data uploads in a database To mitigate sensitive data exposure: Encrypt data in transit and at rest using cryptographic algorithms and keys. Disable caching on forms that collect data. Implement hashed and salted passwords.
Fuzz testing	Fuzz testing (also known as fuzzing) is a software testing technique that exposes security problems by providing invalid, unexpected, or random data to the inputs of an application. Fuzzing program types are: Mutation-based Mutate existing data samples to create data Generation-based Define new test data based on models of the input

Code Signing

Code signing practices use digital signatures to verify the integrity and authenticity of software code. Code signing serves a dual purpose: ensuring that software has not been tampered with since signing and confirming the software publisher's identity.

When software is digitally signed, the signer uses a private key to encrypt a hash or digest of the code—this encrypted hash and the signer's identity form the digital signature. Code signing requires using a certificate issued by a trusted certificate authority (CA). The certificate contains information about the signer's identity and is critical for verifying the digital signature. If the certificate is valid and issued by a trusted CA, the software publisher's identity can be confidently verified. Code signing helps analysts and administrators block untrusted software and also helps protect software publishers by providing a mechanism to validate the authenticity of their code. Overall, code signing helps build trust in the software distribution process.

While code signing provides assurance about the origin of code and verifies code integrity, it does not inherently assure the safety or security of the code itself. Code signing certifies the source and integrity of the code, but it doesn't evaluate the quality or security of the code. The signed code could still contain bugs, vulnerabilities, or malicious code inserted by the original author. Signing ensures software is from the expected developer and in the state the developer intended. While code signing adds trust and authenticity to software distribution, it should not be relied upon to guarantee secure or bug-free code.

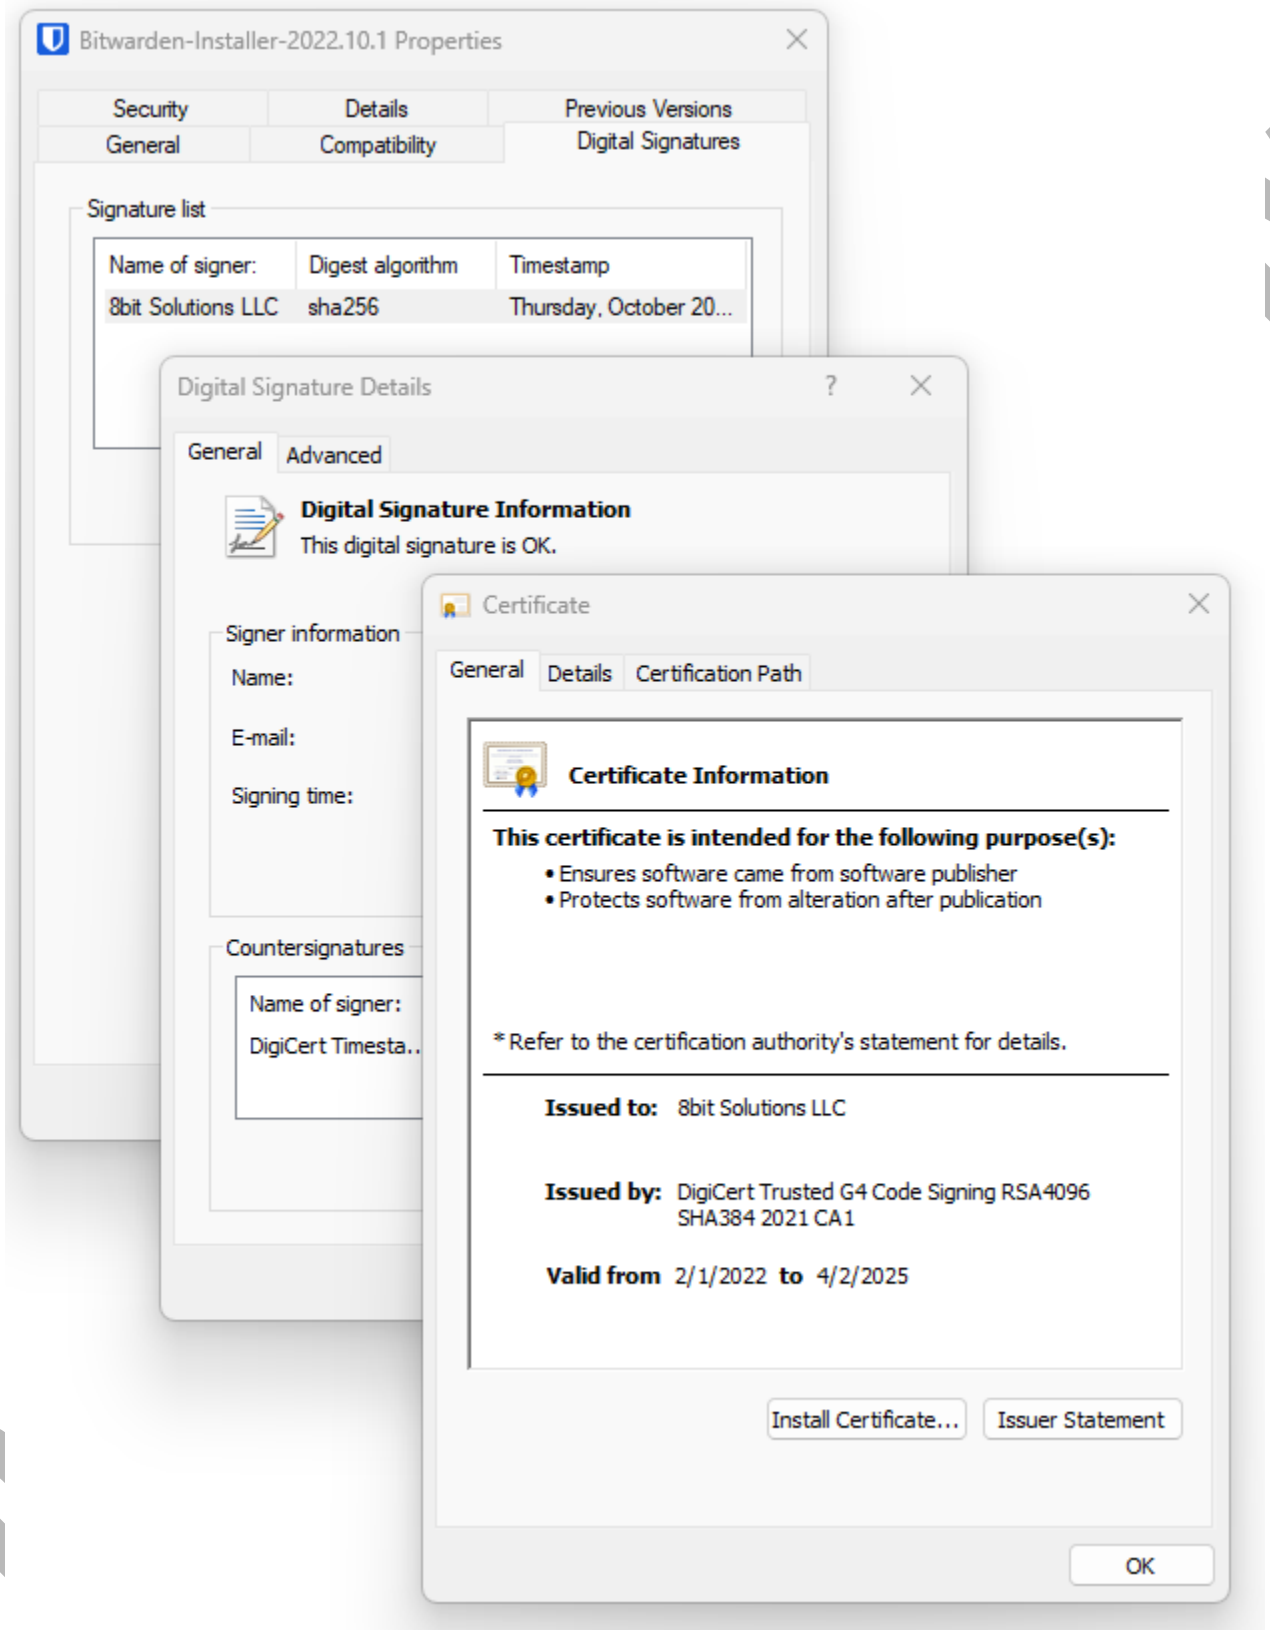

Reviewing the digital s ignature contained within the B itwarden Password Management app installer.

Data exposure is a fault that allows privileged information (such as a token, password, or personal data) to be read without being subject to the appropriate access controls. Applications must only transmit such data between authenticated hosts, using cryptography to protect the session. When incorporating encryption in code, it is important to use industry-standard encryption libraries that are proven to be strong, rather than i nternally developed ones.

Secure Cookies

Cookies are small pieces of data stored on a computer by a web browser while accessing a website. They maintain session states, remember user preferences, and track user behavior and other settings. Cookies can be exploited if not properly secured, leading to attacks such as session hijacking or cross-site scripting.

To implement secure cookies, developers must follow certain well-documented principles, such as using the 'Secure' attribute for all cookies to ensure they are only sent over HTTPS connections and protected from interception via eavesdropping, using the 'HttpOnly' attribute to prevent client-side scripts from accessing cookies and protect against cross-site scripting attacks, and using the 'SameSite' attribute to limit when cookies are sent to mitigate cross-site request forgery attacks. Additionally, cookies should have expiration time limits to restrict their usable life.

Secure cookie techniques are critical in mitigating several web-based application attacks, particularly those focused on unauthorized access or manipulation of session cookies. Developers can defend against attacks that target them by employing specific attributes within cookies.

8.9.6 Hardening Applications on Linux (Demo Video)

Transcript:

In this demonstration, we're going to look at how to harden applications or services on a Linux system. There are several different services you can use to harden applications or services. One of them is called SELinux. It's a cool access control security feature.

Essentially, the idea behind SELinux is that you use the operating system kernel to limit what a particular application is allowed to do, and it goes beyond the limitations that you set with permissions or privileges based on the user account. Think of it as a roadblock to help prevent hackers from doing more damage.
SELinux can be set to determine access between users, files, directories, memory, ports, and sockets. First, we're going to type â€˜sestatus'. this will tell us if SELinux is enforcing the system. Next, we'll take a look at the SELinux config file located in â€˜/etc/selinux/config'. This config file sets the permanent state of SELinux. It's recommended to set your system to permissive first, install all your needed applications, and then view the logs to make the necessary changes to SELinux. The permissive setting logs all the permission denied errors while allowing an application to run. Any changes to the SELinux config file located in /etc/selinux/config requires a reboot.
Two places we can look for SELinux permission denied errors are in the audit log and journalctl. Let's take a look at the audit log, â€˜less /var/log/audit/audit.log'. We're going to use a question mark and search for â€˜AVC', which is a type of SELinux error. As you can see, there's some denied messages from other applications that we should probably look at. Let's exit out of here and look at â€˜journalctl'. SELinux errors may be in red if there's a problem to address. We don't have much time to search through the logs right now, but this should give you a good idea of where to look.
If a service is unable to start, it might cause SELinux to deny permission to certain aspects that allow the application to run. When you need to determine the cause of a problem, logs are your friend. Let's type â€˜clear' to clear our screen. By default, SELinux has hundreds of Booleans that are essentially off and on switches for certain tasks. To view the whole list of Booleans, type â€˜getsebool -a'. Let's say I only want to see Booleans related to samba. Let's type â€˜getsebool -a | grep samba'. To enable a Boolean, we can use the setsebool command. There's one in particular that we want to enable called samba_share_nfs. This one grants samba rights to export NFS volumes. To set this, type â€˜setsebool samba_share_nfs 1 -P'. Now, if we go back and type â€˜getsebool -a | grep samba', you'll see that samba_share_nfs is set to On.
When SELinux is enforced, it uses a categorization system called context types. These are labels with addition information related to its purpose, user, role, and such. Today, we're just going to cover the basics. One useful command you can use to view the context type is the -Z flag. Let take a look at the context types for the samba configuration folder.

Type â€˜ls -lZ /etc/samba'. As you can see, there's additional information listed here that gives a label to samba so SELinux knows these are samba configuration files. If a configuration file was made in a /home directory and then copied to an /etc config directory, it may not have the right context type when copied. This could prevent applications from working correctly and show a denied message from SELinux in the logs.
And that's it for this demonstration. In this demo, we looked at hardening Linux applications using SELinux. We talked about how SELinux can protect your system using access control security. Then we saw some of the commands for managing SELinux and how to troubleshoot permission denied errors from SELinux.

8.9.7 Implementing Application Whitelisting with AppLocker (Demo Video)

Transcript:

In this demonstration, we'll practice implementing application deny and allow lists using AppLocker. AppLocker is a very useful feature of Windows. It allows you to control what applications are allowed or not allowed to run on machines in your domain.

As a real-world example of why you might do this, I had a junior administrator once that ran some registry cleaning software on a production SQL server, and it damaged SQL. The database would still run, but I was not able to upgrade it to a newer version of SQL. After I set up a new SQL server and got the database moved over and functioning, I took the time to set up some AppLocker rules. I downloaded every registry cleaning software I could find and made Deny rules in AppLocker, so no one would ever be able to run a registry cleaner on a production server again.
You can use it to accomplish such things as preventing malware from being installed, preventing unsupported applications from being installed.
The process of implementing application block and allow lists is fairly complex.
First, you must implement the rules for executable files. Then, you test those rules to verify that they are working as you expect without negatively affecting users.
After you analyze events created by AppLocker, you can then modify the rules and enforce them.
Let's begin with the first step of defining the executable rules. In this scenario, we'll enforce the AppLocker policy settings using Group Policy.
By doing this, anytime a user logs into the domain, the AppLocker settings are automatically enforced. Let's open the Group Policy Manager.
The default domain policy would be a bad choice of Group Policy objects (GPOs) to use. Improperly configured Applocker rules can render your system completely inoperable. If you put your rules in the Default Domain Policy and they don't work as expected, you're going to cause widespread problems across your whole domain. So, let's create a new GPO called Applocker. Let's right-click it and select Edit. We're going to resize this screen so we can see what's going on a little better.
Under Computer Configuration, we'll expand Policies > Windows Settings > Security Settings > Application Control Policies.
When we expand AppLocker, we can see these sets of application rules we can use. Let's select AppLocker and click Configure rule enforcement.
For our purposes today, we'll focus on Executable rules. We'll select the Configured checkbox here. After that we need to define whether we want to enforce those rules or just audit them.
We'll choose to audit only for now. You could immediately enforce the rules if you wanted, however, in a production environment you want to make sure that the rules you define work properly for the end users in the organization.
When we select Audit, the events are generated by AppLocker and we can access those in Event Viewer. We can analyze Event Viewer events to decide whether the rules are working correctly or whether we need to modify them.
Let's go to Executable Rules. Notice that there are no rules here. We're enforcing them, but there are no rules to enforce. So, let's right-click here and select Create New Rule.
In the Before You Begin screen, it tells you how the rule works and some preparatory steps you need to follow.
For example, you need to make sure that the applications you want to create the rule for are installed on this system. We've already done that. There are other steps, but we don't need to worry about them. Let's click Next.
Now we have to specify the action to be used in the AppLocker rule. As noted up here, an allow action permits the affected files to run.
A deny action, on the other hand, prevents the files from running. We want to use an Allow rule. Another thing we do on this page is specify who this rule applies to.

The application we are allowing is called SalesBuilder. It's used by the Sales department. As such, we want this rule to apply only to the Sales group.
So, let's specify that this rule applies to the Sales group by clicking Select. We'll type in Sales and click Check Names. Windows has found the group, so we'll click OK. Click Next.
At this point, we need to specify the primary condition that will be used in the rule. We have three options. We'll use the Path option.
We would use the Publisher option if the application we want to create the rule for is signed by a known software publisher, but since ours is an in-house application that isn't signed, we can't' use that option. Let's click Next.
Now we need to specify the path to the application. We'll click Browse Folders and let's browse to the folder where the executable resides. It's SalesBuilder; select it. Notice that we have two options for specifying a path.
We can specify a file or specify a folder. Because we want the rule to apply to all files in the SalesBuilder folder, we specified a folder instead of a file. Click Next. We don't want to create any exceptions. The default rule name isn't very friendly, so I'm going to give this rule a more descriptive name.
Notice that we're presented with this dialogue here. It tells us that we've created a custom rule, but there are default rules that should be created so that important system files will still be able to be run.
If you don't create these default rules, you may not be able to run most of the applications that you need. Let's create the default rules along with the rule we just defined by clicking Yes.
And there you go. Notice that the default rules allow basically full access to anyone in the administrator's group.
The Everyone group is allowed access to all files located in the Windows folder and it is also allowed access to files that are installed in the Program Files folder.
We had to create this custom rule because SalesBuilder is not installed in the Program Files folder. It's installed in the root directory of the C: drive. We created a special rule that allows the Sales group to run the application.
It's important to know that the rules we just configured won't be applied to any users who are currently logged in while we defined these rules.
The rules won't take effect until those users log out and log back in, or until you run GPUpdate on the workstation.
Now, it's important to note that AppLocker uses the Application Identity service on individual workstations to verify the attributes of the executable files we specified in our executable rules. Therefore, we need to make sure that the Application Identity service is always started by default. We can do this using Group Policy.
By doing this when the user logs into the domain, we enforce the AppLocker rules and we make sure that the Application Identity service automatically runs as well.
So, let's edit our AppLocker policy again. We'll go to Computer Configuration > Policies > Windows settings > Security Settings.
Under Security Settings let's select System Services. Over here in the right pane, we'll right-click Application Identity. And then we'll define this to start automatically. Click OK.
With the rules configured and the Application Identity service starting automatically, we can test our AppLocker policy. We can have end users log in to the domain, run the application as they normally would, and then we analyze the AppLocker events.
Based on what we see in Event Viewer, we would decide whether the rules we set up are functioning correctly. If they are, we can enforce the policy. If not, we would go in and modify the policy to work the way it should.
That's it for this demonstration. In this demo, we implemented application block and allow lists with AppLocker.

8.9.8 Implement Application Whitelisting with AppLocker (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You are increasing network security by implementing application whitelisting.

Your first step is to prevent applications not located in the operating system directory or the program files directory from running on your computers. In addition, the call center application used by the support team runs from C:\CallCenter\CallStart.exe and must be allowed to run. You also want any future versions of the call center application to run without changing any settings.

In this lab, your task is to configure AppLocker in the default domain policy as follows:

- Create the default rules.

Allow all files located in the Program Files folder.

Allow all files located in the Windows folder.

Configure a publisher rule that will allow future updates from the same vendor.

Allow the **Support** group to run the call center software found in **C:\CallCenter\CallStart.exe** .

Explanation

Complete this lab as follows:

Access the CorpNet.local domain under Group Policy Management.

From the Server Manager's menu bar, select **Tools** > **Group Policy Management** .

Maximize the window for better viewing.

Expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** .

Access the AppLocker policy.

Right-click **Default Domain Policy** and select **Edit** .

Maximize the window for better viewing.

Under *Computer Configuration* , expand and select:
Policies > **Windows Settings** > **Security Settings** > **Application Control Policies** > **AppLocker** .

Configure rule enforcement.

From the right pane, select **Configure rule enforcement** .

Under Executable rules, select **Configured** .

Make sure **Enforce rules** is selected in the drop-down list.

Select **OK** .

Configure a Publisher rule and allow the **Support** group to run the call center software.

From the left pane, expand **AppLocker** .

Right-click **Executable Rules** and then select **Create New Rule** .

Select **Next** .

Make sure **Allow** is selected.

For *User or group* , click **Select** .

In the *Enter the object names to select* box, type **Support** , and then select **OK.**

Select **Next** .

Make sure **Publisher** is selected, and then select **Next** .

For the Reference file, select **Browse** .

Browse to and select the **C:\CallCenter\CallStart.exe** file.

Select **Open** .

Slide the pointer from *File version* to **Publisher** and then select **Next** .

Select **Next** .

Accept the ***default name*** and select **Create** .

Select **Yes** to create the default rules.
Notice that the Publisher rule was created.

8.9.9 Implementing Data Execution Preventions (Demo Video)

Transcript:

In this demonstration, we're going to practice working with Data Execution Prevention, or DEP.

DEP is a set of hardware and software technologies designed to prevent malicious code from running on a Windows system.
DEP looks for and blocks a malicious program or any other type of attack that has injected a process with malicious code and is trying to run that code. Let's look at how this works.
Let's start with a Windows system. If I right-click Start, select System, come over to Advanced system settings, go to Performance Settings on the Advanced tab, and then go to the Data Execution Prevention tab. In Windows Server 2022, DEP is enabled by default for all programs, except those the user selects.
It's possible that DEP will trigger a false positive. This means that DEP thinks a legitimate process that's running on the system is malicious and shuts it down.
This is particularly relevant in situations where you have in-house programmers who create custom applications for your organization.
In-house developed applications sometimes aren't as rigorously tested, so they might kick off an exception where DEP detects them and thinks that they're malicious code.
But notice you do have the option to turn off DEP for specific applications. We have a case where we can do that on this system.
I'm going to come down here to File Explorer. If I go to my C: drive, then into Program Files, and then to the SalesBuilder folder, there's an application called SalesBuilder.
This is an in-house written application. It wasn't extensively tested and will kick off an exception when we try to run it if DEP is monitoring all system processes.
Let's hit Add and go to Program Files > SalesBuilder > SalesBuilder.exe. Now DEP is turned on for everything except SalesBuilder. We hit Apply.
Verify that DEP is Enforced
DEP is now enforced. DEP is now on for all services except SalesBuilder.
This means that whenever a process is run on this system, whether it's a service or application, DEP will watch it and look for any instance where it's trying to run code from places that it shouldn't. Except, of course, for the one application that I specified earlier, SalesBuilder, which DEP was told not to worry about because I know that I trust it.

That's it for this demonstration. In this demo, we configured Data Execution Prevention, or DEP.

8.9.10 Implement Data Execution Preventions (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. You are configuring the computer in Office 1 to use Data Execution Prevention (DEP) for all programs and services. You have noticed that the accounting program used on some computers does not function well when DEP is enabled.

In this lab, your task is to configure DEP as follows:

Enable DEP for all files.

Disable DEP for **C:\Program Files (x86)\AccountWizard\AccountWizard.exe** .

Restart the computer to activate DEP.

Explanation

Complete this lab as follows:

Access the Advanced system settings (System Properties).

Right-click **Start** and then select **System** .

Maximize the window for better viewing.

From the right pane, under *Related settings* , select **System protection** .

Select the **Advanced** tab.

Configure Data Execution Prevention.

Under Performance, select **Settings** .

Select the **Data Execution Prevention** tab.

Select **Turn on DEP for all programs and services except those I select** .

Select **Add** .

Open the **C:\Program Files (x86)\AccountWizard** folder.

Select **AccountWizard.exe** .

Select **Open** .

Make sure **AccountWizard.exe** is selected, and then click **OK** .

Select **OK** to confirm that a system restart is needed.

Select **OK** to close System Properties.

Select **Restart Now** to restart the computer and activate DEP.

Wait for the computer to restart.

8.9.11 Hardening Applications Facts

Application hardening is the process of preventing the exploitation of vulnerabilities in software applications. Applications pose the most difficult security challenges for a security administrator because they are complex, usually developed by a third party, and designed to accept input from users.

This lesson covers the following topics:

Application hardening guidelines

Application hardening techniques

Input validation

Monitoring capabilities

Application Hardening Guidelines

Basic hardening guidelines for applications are as follows:

Assume all installed applications are flawed.

Remove all unused applications from the system.

Limit administrative privileges.

Install security software, such as antivirus, anti-spyware, and anti-rootkit software.

Use firewalls, content filters, and operating system user-authentication features.

Restrict access to the application and provide access only to those who need it.

Update all applications with the latest patches when security bulletins are released.

Identify baselines.

Application Hardening Techniques

Additional application hardening includes the following techniques:

Technique	Description

Block process spawning	Process spawning is the creation of a new process (also called a child process) by an existing process (also called a parent process). If you take the process-spawning ability from the application, threat agents cannot perform process-spawning attacks.
Control access to executable files	Executable files should be protected from modification by removing the Write permissions given to applications.
Protect OS components	Sensitive file system areas (such as Windows Registry keys) should be protected by removing Write permissions given to specific applications. In most cases, applications do not need to modify sensitive areas of the system for them to function properly.
Use exception rules	Exception rules allow an administrator to bypass a specific hardening rule when an application has a legitimate need. Administrators should exercise caution and set parameters regarding the exceptions that ensure the security of the system.
Use Data Execution Prevention	Data Execution Prevention (DEP) is a security feature that can help prevent damage to your computer from viruses and other security threats. DEP ensures that applications use computer memory safely. DEP closes an application and notifies the administrator if a program initiates run instructions from the portion of memory used for data.
Implement third-party application hardening tools	Third-party application hardening systems are developed for specific applications. The rules used by the application hardening system can be applied to the application being hardened, including libraries and SDKs. An example of these tools is AppArmor for Linux systems.

Input Validation

Input validation is an essential protection technique used in software and web development that addresses the issue of untrusted input. Untrusted input describes how an attacker can provide specially crafted data to an application to manipulate its behavior. Injection attacks exploit the input mechanisms applications rely on to execute malicious commands and scripts to access sensitive data, control the operation of the application, gain access to otherwise protected back-end systems, and disrupt operations.

OWASP provides an excellent overview of input validation at https://cheatsheetseries.owasp.org/cheatsheets/Input_Validation_Cheat_Sheet.html .

Without effective input validation, applications are vulnerable to many different classes of injection attacks, such as SQL injection, code injection, cross-site scripting (XSS), and many others.

Monitoring Capabilities

Secure coding practices focus primarily on preventing software vulnerabilities and stress enhancements to logging and monitoring capabilities. These features support security analysts tasked with detecting potential threats and malicious activity in software. Writing code with enhanced monitoring capabilities improves the granularity and effectiveness of logging and alerting systems, which are crucial system monitoring tools.

Implementing comprehensive and meaningful logging requires developers to ensure their applications generate logs that capture important events and activities to support security audits, incident response, and system troubleshooting. Secure coding practices encourage robust error handling to hide or mask sensitive debugging information, and this practice minimizes the risk of attackers

exploiting information displayed in error messages. Integrating real-time alerting capabilities within the application code can significantly improve threat detection. For example, code that triggers alerts when specific events occur, such as repeated failed login attempts or unusual data transfers, helps security analysts monitor applications more effectively. These alerts often indicate a potential security breach and provide crucial information for incident response teams.

8.9.12 Practice Questions (Section Quiz)

q_sdlc_agile_secp8

Which application development model approaches software development as a continuous, changing process with never-ending versions, bug fixes, and enhancements?

Answers:

***Agile**

Waterfall

Fuzz testing

Code signing

Explanation:

The Agile development model approaches software development as a continuous, changing process with never-ending versions, bug fixes, and enhancements.

The Waterfall development model is the most widely used model. It is called this because each step is completed before the next step is begun. This way, each step flows to the next.

Fuzz testing is a software testing technique that exposes security problems by providing invalid, unexpected, or random data to the inputs of an application.

Code signing is the process of digitally signing (encrypting) executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed.

q_sdlc_catchall_handler_secp8

You are a software developer working on a complex application. During the testing phase, you notice that the application crashes unexpectedly when certain unpredictable events occur.

You have implemented a try-catch block to handle anticipated errors, but the application still crashes on unanticipated errors.

What should you do to ensure the application handles these errors gracefully and does not reveal any sensitive information about the system or the code?

Answers:

***Implement a catchall handler.**

Ignore the unanticipated errors, as they are rare.

Use default error handlers provided by the application's interpreter.

Remove the try-catch block, as it is not effective.

Explanation:

Implementing a catchall handler is the correct solution. A catchall handler is designed to handle all types of errors, including those that are unanticipated. This ensures that the application does not crash and does not reveal sensitive information when an unexpected error occurs.

Ignoring the unanticipated errors is not a good solution. Even though these errors might be rare, they can still cause the application to crash or behave unpredictably. Ignoring them does not solve the problem and can lead to serious security issues.

Using default error handlers provided by the application's interpreter is not the best solution. These handlers often display default error messages that can reveal platform information and the inner workings of the code to an attacker.

Removing the try-catch block is not a good solution. The try-catch block is designed to handle anticipated errors. Removing it would leave these errors unhandled, which could lead to crashes or unpredictable behavior.

q_sdlc_dast_secp8

A software development team is working on an application that will handle sensitive user data. The team has already implemented static application security testing (SAST) in the early stages of development.

As the application nears deployment, what additional secure testing method should the team consider implementing and why?

Answers:

***Dynamic application security testing (DAST) because it can test the application after deployment and from the outside.**

Additional static application security testing (SAST) because it can identify the exact cause of a coding problem.

Interactive application security testing (IAST) because it can access interpreters and compilers, allowing precise identification of a problematic line of code in runtime.

No additional testing is needed because SAST is sufficient for ensuring the security of the application.

Explanation:

Dynamic application security testing (DAST) because it can test the application after deployment and from the outside is the correct answer. Dynamic application security testing (DAST) is a good complement to SAST because it tests the application in its running state from the outside, mimicking an attacker's perspective. This can help identify vulnerabilities that SAST might miss, such as runtime errors and server configuration issues.

While additional SAST can be beneficial, it is not the best choice in this scenario. SAST is most effective in the early stages of development and primarily focuses on analyzing source code, binaries, and byte code. It may not identify vulnerabilities that only become apparent when the application is running.

Interactive application security testing (IAST) can be a good choice, but it is not the best in this scenario. While IAST can identify problematic lines of code in runtime, it requires a higher level of integration with the application and may not be as effective as DAST in identifying vulnerabilities from an external perspective.

While SAST is an important part of secure testing, it should not be the only method used. SAST focuses on the internal structure of the application and may not identify vulnerabilities that only become apparent when the application is running or when it is attacked from the outside.

q_sdlc_dynamic_secp8

You are performing a security test from the outside on a new application that has been deployed.

Which secure testing method are you MOST likely using?

Answers:

Static

Interactive

***Dynamic**

Runtime

Explanation:

Dynamic application security testing scans applications after they have been deployed. These tests are performed from the outside.

Static application security testing focuses on analyzing source code, binaries, and byte code early in the development process.

Interactive application security testing is built into static testing and uses source code scanners.

Runtime is a type of coding error that occurs while software is running.

q_sdlc_require_secp8

Which of the following is the first step in the Waterfall application development model?

Answers:

Design

Implementation

***Requirements**

Maintenance

Explanation:

The Waterfall development life cycle model steps are:

Requirements

Design

Implementation

Testing

Development

Maintenance

q_sdlc_seh_secp8

You are a software developer working on a new application. During the testing phase, you notice that the application crashes unexpectedly when certain unpredictable events occur, such as loss of network connectivity or invalid user input.

You need to ensure that the application handles these errors gracefully and does not reveal any sensitive information about the system or the code.

Which of the following would be the MOST effective solution to this problem?

Answers:

***Implement a structured exception handler (SEH).**

Use default error handlers provided by the application's interpreter.

Ignore the errors as they are unpredictable and cannot be prevented.

Use a try-catch block without a catchall handler.

Explanation:

Implementing a structured exception handler (SEH) is the correct solution. SEH is a mechanism that allows a program to intercept and respond to exceptional conditions such as errors. It allows the application to handle errors gracefully and prevents the application from crashing or revealing sensitive information.

Using default error handlers provided by the application's interpreter is not the best solution. These handlers often display default error messages that can reveal platform information and the inner workings of the code to an attacker.

Ignoring the errors is not a good solution. Even though these errors are unpredictable, they can still cause the application to crash or behave unpredictably. Ignoring them does not solve the problem and can lead to serious security issues.

Using a try-catch block without a catchall handler is not the best solution. While this approach can handle some anticipated errors, it does not provide a solution for unexpected errors. Without a catchall handler, unexpected errors can cause the application to crash or behave unpredictably.

q_sdlc_software_sandboxing_secp8

You are a security analyst at a software development company. Your company is developing a new application that will handle sensitive user data.

To ensure the security of the application, you are considering various security measures. One of your colleagues suggests using software sandboxing.

How would implementing software sandboxing contribute to the security of the application?

Answers:

It would encrypt the user data to prevent unauthorized access.

***It would isolate the application from the rest of the system, limiting the potential damage from security vulnerabilities.**

It would detect and remove malware from the application.

It would detect and remove malware from the application.

Explanation:

Isolating the application from the rest of the system is the correct answer. Software sandboxing is a security mechanism that isolates an application in a separate environment to prevent it from interacting with the rest of the system. This limits the potential damage that can be caused if the application has security vulnerabilities.

Software sandboxing does not directly involve encryption. While encryption is a valuable security measure for protecting data, it is not the primary function of a sandbox.

While sandboxing can help to contain malware if it is present within the sandboxed application, it does not actively detect or remove malware. Other security measures, such as antivirus software, are needed for this purpose.

Software sandboxing does not prevent users from inputting invalid data. Input validation is a separate security measure that needs to be implemented in the application's code.

q_sdlc_software_secp8

Which of the following are the two main causes of software vulnerabilities? (Select two.)

Answers:

***Coding errors**

Fuzzing

Normalization

***Design flaws**

Obfuscation

Explanation:

Coding errors and design flaws are the main causes of software vulnerabilities.

Fuzz testing (also known as fuzzing) is a software-testing technique that exposes security problems by providing invalid, unexpected, or random data to the inputs of an application.

Normalization is data reorganized in a relational database with the intent to eliminate redundancy. This is done by having all related data stored in one place. This is not one of the main causes of software vulnerabilities.

Obfuscation is the deliberate act of creating source or machine code that is difficult for humans to understand. This is not one of the main causes of software vulnerabilities.

q_sdlc_sprint_secp8

A software development team is using the Agile development life cycle model for a new project.

During a sprint, they realize that a feature they are working on is more complex than initially thought and cannot be completed within the current sprint.

What should the team do?

Answers:

Continue working on the feature past the end of the sprint until it is completed.

Drop the feature from the project.

***Break the feature down into smaller tasks and distribute them over the next few sprints.**

Move on to the next feature and return to the complex feature in a future sprint.

Explanation:

Breaking the feature down into smaller tasks and distributing them over the next few sprints is the correct answer. In Agile, if a task or feature is too large to be completed within a sprint, it should be broken down into smaller, manageable tasks. These tasks can then be distributed over the next few sprints. This approach maintains the time-boxing principle of Agile and allows for complex features to be developed in a manageable way.

One of the principles of Agile is time-boxing, where work is confined to a regular, repeatable work cycle. Extending the sprint to accommodate a complex feature goes against this principle and can disrupt the project timeline.

Dropping the feature from the project is not a good solution unless the feature is deemed unnecessary or unfeasible. Agile promotes adaptive planning and evolutionary development, but it doesn't mean dropping a feature just because it's complex.

Moving on to the next feature and returning to the complex feature in a future sprint is not the best choice. While it's important to keep the project moving, simply postponing the complex feature may lead to a backlog of unfinished tasks. It's better to break down the feature into smaller tasks and tackle them over the next few sprints.

q_sdlc_waterfall_secp8

Which of the following is considered a drawback of the Waterfall application development life cycle?

Answers:

Development is broken into Sprints.

Testing is performed throughout development.

***Requirements are determined at the beginning and are carried through to the end product.**

Each step in the life cycle only needs to be completed once before moving on to the next one.

Explanation:

The Waterfall development life cycle is a slow process and may take months or years to complete. It also lacks flexibility since the requirements determined in the beginning are carried through to the end product.

Development is broken into Sprints when using the Agile development model.

The Agile development model performs testing throughout development.

When using the Waterfall development model, an application likely goes through some of these steps multiple times before moving on to the next step.

q_app_devsec_code_obfuscation_secp8

John, a software developer, is working on a project that involves creating a proprietary algorithm for his company.

He wants to ensure that even if someone gets access to the source code, they would have a hard time understanding the logic and purpose of the code.

Which secure coding concept should John use?

Answers:

Code reuse

***Code obfuscation**

Stored procedures

Normalization

Explanation:

Code obfuscation is the correct answer. Code obfuscation involves making the source or machine code difficult for humans to understand. It can be used to conceal the purpose or logic of the code, making it a good choice for John's situation.

Code reuse involves using the same code multiple times. While it can increase efficiency, it does not inherently protect the code from being understood if accessed.

Stored procedures are a group of one or more database statements stored in a database's data dictionary. While they can centralize code and limit injection attacks, they do not make the code difficult to understand.

Normalization is a process in database design to minimize redundancy and dependency by organizing fields and tables of a database. It does not make the code difficult to understand.

q_app_devsec_code_secp8

You have just finished developing a new application. Before putting it on the website for users to download, you want to provide a checksum to verify that the object has not been modified.

Which of the following would you implement?

Answers:

Memory management

Code obfuscation

***Code signing**

Normalization

Explanation:

Code signing is the process of digitally signing (encrypting) executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed. The process employs the use of a cryptographic hash to validate authenticity and integrity.

Code signing:

Provides security when deployed.

Helps prevent namespace conflicts in some programming languages.

Provides a digital signature mechanism to verify the identity of the author or build system.

Provides a checksum to verify that the object has not been modified.

Provides versioning information about an object as well as storing other metadata about the object.

Memory management is a resource-management process applied to computer memory.

Code obfuscation is the deliberate act of creating source or machine code that is difficult for humans to understand.

Normalization is data reorganized in a relational database with the intent to eliminate redundancy by having all related data stored in one place.

q_app_devsec_fuzz_secp8

Which fuzz testing program type defines new test data based on models of the input?

Answers:

***Generation-based**

Code signing

Memory management

Mutation-based

Explanation:

Fuzz testing (also known as fuzzing) is a software-testing technique that exposes security problems by providing invalid, unexpected, or random data to the inputs of an application. Fuzzing program types are:

Mutation-based programs

Mutate existing data samples to create data

Generation-based programs

Define new test data based on models of the input

Code signing is the process of digitally signing (encrypting) executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed.

Memory management is a resource-management process applied to computer memory.

q_app_devsec_memory_management_secp8

You are a software engineer working on a large-scale application.

During a code review, you notice that a fellow developer has written a function that allocates a large block of memory but does not deallocate it before the function ends.

What potential issue could this lead to?

Answers:

Code obfuscation

***Memory leak**

Code reuse

Fuzz testing

Explanation:

Memory leak is the correct answer. A memory leak occurs when a program allocates memory but fails to deallocate it, leading to a gradual loss of available memory. This can eventually cause the program or even the entire system to crash due to lack of memory.

Code obfuscation involves making the source or machine code difficult for humans to understand. It does not directly relate to memory management or the allocation and deallocation of memory.

Code reuse involves using the same code multiple times. While it can increase efficiency, it does not directly relate to memory management or the allocation and deallocation of memory.

Fuzz testing is a software testing technique that involves providing invalid, unexpected, or random data to the inputs of an application. While it can help identify security vulnerabilities, it does not directly address the issue of memory management.

q_app_devsec_sdk_secp8

What is a set of software development tools called that can be installed as one unit and provides code frameworks or code snippets to help development go faster?

Answers:

Repository

***SDK**

Code signing

Memory management

Explanation:

A software development kit (SDK) is a set of software development tools that can be installed as one unit. These tools can provide code frameworks or code snippets to help development go faster.

A version control system uses a repository, which is a storage location that holds all the source files used during development.

Code signing is the process of digitally signing (encrypting) executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed.

Memory management is a resource-management process applied to computer memory.

q_app_devsec_secure_cookies_secp8

You are a security analyst at a tech company developing a new web application. The development team has implemented cookies to maintain user sessions. You are tasked with ensuring the security of these cookies.

Which of the following approaches would you recommend to the team?

Answers:

Implement secure cookies but allow them to be transmitted over HTTP connections.

***Implement secure cookies and ensure they are only transmitted over HTTPS connections.**

Avoid using cookies altogether and use URL parameters to maintain user sessions.

Implement secure cookies and store user credentials in them for easy access.

Explanation:

Implementing secure cookies and ensuring they are only transmitted over HTTPS connections is the most secure approach. HTTPS encrypts the data between the client and the server, which helps protect the cookies from being intercepted and read by unauthorized parties. Therefore, this is the correct answer.

Implementing secure cookies but allowing them to be transmitted over HTTP connections is not a secure approach. Even though the cookies are secure, transmitting them over HTTP can expose them to man-in-the-middle attacks. Therefore, this is not the best answer.

Avoiding the use of cookies altogether and using URL parameters to maintain user sessions is not a secure approach. URL parameters are visible in the browser's address bar and can be stored in browser history, bookmarks, and web server logs. This could potentially expose sensitive session information. Therefore, this is not a good answer.

Implementing secure cookies and storing user credentials in them for easy access is a very insecure approach. If an attacker manages to steal the cookie, they would have direct access to the user's credentials. Therefore, this is not a good answer.

q_app_devsec_sensitive_data_secp8

As a lead software developer, you are reviewing the code of a new application developed by your team. The application collects personal information from users, including their names, addresses, and credit card details.

You notice that the data is stored in plain text in the database and transmitted over the network without any encryption.

What security issue is MOST directly associated with this situation?

Answers:

Code obfuscation

Code reuse

***Sensitive data exposure**

Fuzz testing

Explanation:

Sensitive data exposure is the correct answer. Sensitive data exposure involves unintended exposure of personal and confidential data. Storing and transmitting sensitive data in plain text without any encryption is a direct example of sensitive data exposure.

Code obfuscation involves making the source or machine code difficult for humans to understand. While it can be used to protect the logic and purpose of the code, it does not directly address the issue of protecting sensitive data.

Code reuse involves using the same code multiple times. While it can increase efficiency, it does not inherently protect sensitive data.

Fuzz testing is a software testing technique that involves providing invalid, unexpected, or random data to the inputs of an application. While it can help identify security vulnerabilities, it does not directly address the issue of protecting sensitive data.

q_app_hard_exception_rules_secp8

As a security administrator, you have been tasked with hardening a critical application in your organization's IT infrastructure.

The application is complex and developed by a third party. It is also designed to accept input from users. The application has unique functionalities necessary for business operations but may pose security risks.

Which of the following application hardening techniques would be the BEST solution for you to use?

Answers:

Block process spawning.

Control access to executable files.

***Use exception rules.**

Implement third-party application hardening tools.

Explanation:

Exception rules allow an administrator to bypass a specific hardening rule when an application has a legitimate need. This technique is important for maintaining the functionality of the application while also ensuring its security. Given the unique functionalities of the application that are necessary for business operations, using exception rules would be the BEST solution for you to use.

Block process spawning involves preventing the creation of a new process (also called a child process) by an existing process (also called a parent process). Blocking process spawning can help prevent process spawning attacks. However, this might not be the highest priority if the application does not have functionality that involves process spawning or if the unique functionalities of the application require process spawning.

Control access to executable files involves protecting executable files from modification by removing the Write permissions given to applications. This is a critical step in application hardening as it prevents unauthorized changes to the application's executable files. However, this might not be the highest priority if the application's executable files are already well-protected and the unique functionalities of the application require certain permissions.

Third-party application hardening tools are developed for specific applications and can be applied to the application being hardened, including libraries and SDKs. While this could be a comprehensive approach to hardening the application, it might not be the highest priority if the unique functionalities of the application require specific exception rules that may not be covered by third-party tools.

q_app_hard_execution_prevention_secp8

You are a security administrator for a large corporation.

You've noticed that one of your applications is being exploited by a threat agent who initiates run instructions from the portion of memory used for data, causing system instability.

Which application hardening technique would be the MOST effective in this situation?

Answers:

Use exception rules

Monitor logs

***Use Data Execution Prevention**

Block process spawning

Explanation:

Use Data Execution Prevention is the correct answer. Data Execution Prevention (DEP) is a security feature that can help prevent damage from viruses and other security threats. DEP ensures that applications use computer memory safely. DEP closes an application and notifies the administrator if a program initiates run instructions from the portion of memory used for data, directly addressing the issue at hand.

Exception rules allow an administrator to bypass a specific hardening rule when an application has a legitimate need. However, in this case, the issue is not about exceptions to hardening rules but about preventing a threat agent from exploiting the application's memory usage.

Monitoring logs is an important part of maintaining system security, as it allows administrators to identify potentially vulnerable applications and detect if an application is being exploited. However, while this technique can help identify the problem, it does not directly prevent the threat agent from exploiting the application's memory usage.

Blocking process spawning is a technique used to prevent the creation of new processes by existing ones. This is not the most effective solution in this case, as the issue is not related to process spawning but to the application's memory usage.

q_app_hard_guideline_secp8

Which of the following is a basic hardening guideline for applications?

Answers:

Grant administrative privileges to all users.

Assume all installed applications are flawless.

***Update all applications with the latest patches when security bulletins are released.**

Allow unrestricted access to the application.

Explanation:

Updating all applications with the latest patches when security bulletins are released is a correct guideline. Regularly updating applications with the latest patches is crucial for maintaining security, as patches often fix known vulnerabilities.

Granting administrative privileges to all users is a security risk. It's better to limit administrative privileges to only those who absolutely need them.

It's a good security practice to assume that all installed applications might have vulnerabilities, as this encourages proactive security measures.

Restricting access to applications helps to minimize the potential attack surface and limit the potential damage if a breach does occur.

q_app_hard_input_validation_secp8

You are a software developer at a tech company and are currently working on a new web application. You are tasked with implementing input validation to prevent potential security threats.

Which of the following approaches would be the MOST effective way to ensure input validation?

Answers:

Implement client-side validation only.

Implement server-side validation only.

***Implement both client-side and server-side validation.**

There is no need for input validation as the application uses secure coding practices.

Explanation:

Implement both client-side and server-side validation is the correct answer. This is the most secure and user-friendly approach. Client-side validation provides immediate feedback to the user, while server-side validation ensures that even if the client-side validation is bypassed, invalid data will not be processed.

While client-side validation can provide immediate feedback to the user and help reduce server load, it can be easily bypassed by a malicious user. Therefore, relying solely on client-side validation is not a secure approach.

Server-side validation is crucial as it cannot be bypassed by the user. However, relying solely on server-side validation may not provide the best user experience as it requires a round trip to the server before the user gets feedback.

Even with secure coding practices, input validation is still necessary. Secure coding practices can help prevent vulnerabilities, but they cannot guarantee that all input will be safe. Input validation is a crucial part of secure coding practices.

q_app_hard_monitoring_secp8

You are a security analyst at a tech company reviewing the monitoring capabilities of a recently developed web application. You notice that the application logs are not capturing some important events and activities.

Which of the following actions would be the MOST effective way to enhance the monitoring capabilities of the application?

Answers:

Increase the frequency of log reviews.

Implement real-time alerting for specific events.

Remove unnecessary log entries to reduce noise.

***Implement a third-party monitoring tool.**

Explanation:

Implementing a third-party monitoring tool is the correct answer. A third-party monitoring tool can provide additional monitoring capabilities that the application may not currently have. This can help ensure that all important events and activities are captured in the logs. However, it's also important to work with the development team to enhance the application's built-in logging capabilities.

While increasing the frequency of log reviews can help identify issues more quickly, it does not address the problem of the application not capturing important events and activities.

Real-time alerting can help detect potential security threats more quickly. However, if the application is not logging important events and activities, real-time alerting may not be fully effective.

While reducing noise in the logs can make it easier to identify important events, it does not address the problem of the application not capturing these events in the first place.

q_app_hard_os_components_secp8

As a security administrator for your company, you have noticed that certain applications have been modifying sensitive areas of the system, leading to potential security vulnerabilities.

Which application hardening technique would be the MOST effective in this situation?

Answers:

Use Data Execution Prevention

***Protect OS components**

Use exception rule

Monitor logs

Explanation:

Protect OS components is the correct answer. Protecting OS components involves removing Write permissions given to specific applications for sensitive file system areas (such as Windows Registry keys). This directly addresses the issue of applications modifying sensitive areas of the system.

While Data Execution Prevention (DEP) is a valuable security feature that can help prevent damage from viruses and other security threats, it is not the most effective solution in this case. DEP ensures that applications use computer memory safely, but it does not directly prevent applications from modifying sensitive system areas.

Exception rules allow an administrator to bypass a specific hardening rule when an application has a legitimate need. However, in this case, the issue is not about exceptions to hardening rules but about preventing applications from modifying sensitive system areas.

Monitoring logs is an important part of maintaining system security, as it allows administrators to identify potentially vulnerable applications and detect if an application is being exploited. However, while this technique can help identify the problem, it does not directly prevent applications from modifying sensitive system areas.

q_app_hard_process_spawning_secp8

You are a security administrator for a large corporation. One of your applications has been experiencing frequent attacks where threat agents are creating new processes from existing ones, leading to system instability and potential data breaches.

Which application hardening technique would be the MOST effective in this situation?

Answers:

Use exception rules.

Monitor logs.

***Block process spawning.**

Implement third-party application hardening tools.

Explanation:

Block process spawning is the correct answer. Blocking process spawning is the creation of a new process (also called a child process) by an existing process (also called a parent process). By taking the process spawning ability from the application, threat agents are not able to perform process spawning attacks, directly addressing the issue at hand.

While exception rules can be useful in certain situations, they are not the most effective solution in this case. Exception rules allow an administrator to bypass a specific hardening rule when an application has a legitimate need, but they do not directly address the issue of process spawning.

Monitoring logs is a crucial part of maintaining system security, as it allows administrators to identify potentially vulnerable applications and detect if an application is being exploited. However, while this technique can help identify the problem, it does not directly prevent process spawning.

While third-party application hardening tools can be useful, they may not specifically address the issue of process spawning. These tools are typically designed to apply a set of predefined rules to an application, and while they may include rules related to process spawning, this is not guaranteed.

9.0 Incident Response

9.1 Incident Response and Mitigation

As you study this section, answer the following questions:

Why is the chain of custody so important in a forensic investigation?

How do you ensure the integrity of collected digital evidence?

When conducting a forensic investigation, what methods can you use to save the contents of memory?

What would a computer forensic investigator analyze when conducting a live analysis compared to a dead analysis?

What actions should you take when an incident occurs?

Key terms for this section include the following:

Term	Definition
Security incident	An event, or series of events, resulting from a security policy violation. A security incident has adverse effects on a company's ability to proceed with normal business.
Incident response	The action taken to deal with an incident, both during and after the incident.
First responder	The first person on the scene after a security incident has occurred.
Damage assessment	A preliminary onsite evaluation of damage or loss caused by a security incident.
Live analysis	An incident investigation that examines an active (running) computer system to analyze the live network connection, memory contents, and running programs.
Dead analysis	An incident investigation that examines data at rest, such as analyzing hard drive contents.
Big data analysis	An incident investigation that examines all types of data used in the organization, including text, audio, video, and log files. The investigation identifies anomalies that led up to the security incident.
Corroborative evidence	Evidence or information that supports another fact or detail.
Hearsay evidence	Evidence that is obtained from a source who doesn't have personal, firsthand knowledge.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Segmentation Isolation 4.8 Explain appropriate incident response activities. Process Preparation Detection Analysis Containment Eradication Recovery Lessons learned 5.1 Summarize elements of effective security governance. Procedures Playbooks
TestOut Security Pro	5.2 Assessment techniques 5.2.4 Analyze network attacks 5.2.5 Analyze password attacks

9.1.1 Incident Response Process (Lesson Video)

Transcript:

Incident response is a systematic approach to handling and managing the aftermath of a security breach or cyberattack, also known as an IT incident, computer incident, or security incident. The goal is to handle the situation in a way that limits damage and reduces recovery time and costs.

The incident response lifecycle provides a structured approach to addressing and managing the aftermath of a security breach or cyberattack to limit damage and reduce recovery time and costs. The incident response lifecycle includes seven steps: preparation, detection, analysis, containment, eradication, recovery, and lessons learned.
Part of the preparation step is to ensure that systems are resilient to attack. This includes hardening systems, writing policies and procedures, and setting up confidential lines of communication.

The preparation step also includes creating incident response resources and procedures. An Incident Response Plan is like a fire drill for cyber threats. Just as a fire drill outlines steps to evacuate a building safely, an Incident Response Plan contains instructions on how to react when a cyber-attack occurs.
The detection step discovers indicators of threat actor activity. Indicators that an incident may have occurred could be generated from an automated intrusion system or other monitoring and alerting systems. Incidents can also be detected using threat hunting methods or by reports made by employees or customers.
The containment phase is where immediate action is taken to prevent further damage or compromise of the system. This could involve disconnecting affected systems or devices from the network to prevent the spread of the breach. It's the equivalent of stopping the bleeding in a medical emergency, a quick and temporary fix to halt the immediate threat. Containment strategies can vary based on the severity and nature of the incident. After containment, an in-depth investigation leads to the next phase—eradication.
Next is the recovery phase. During this phase, systems and networks are returned to their normal function. All systems affected by the cyberattack are cleaned, restored, and put back into operation. Recovery may involve reinstalling system components, changing passwords, and patching software. System administrators should carefully monitor systems during recovery for any signs of abnormal activity, as this could indicate that not all threat elements have been successfully eradicated. Regular operations can resume once the systems are deemed secure and functioning normally. The recovery phase isn't considered complete until all systems are back operational and all data has been recovered.
The last phase in the incident response lifecycle is the "lessons learned" stage. During this phase, the incident response team conducts a post-incident review. This review aims to identify what went well during the response, what could've been done better, and what improvements can be made to the incident response process. It provides an opportunity to learn from the incident and improve future response efforts. The lessons learned may include changes to policies, procedures, or infrastructure. They may lead to further security and awareness training for employees.
That's it for this lesson. In this lesson, we've discussed incident response. We looked at the seven steps of the incident response lifecycle: preparation, detection, analysis, containment, eradication, recovery, and lessons learned.

9.1.2 Incident Response Process Facts

This lesson covers the following topics:

Security incident

Incident response process

Security Incident

A security incident is an event or series of events resulting from a security policy violation. The incident may or may not adversely affect an organization's ability to conduct business. It is crucial to organizations that security incidents are recognized and dealt with appropriately. The following table describes types of security incidents.

Type	Description
Employee errors	Unintentional actions by an employee that cause damage or leave network systems vulnerable to attack.
Unauthorized act by an employee	Intentional actions by an employee to cause harm to a company's network or data. This is also known as an insider threat.
External intrusion attempts	Intentional actions by a threat actor not employed by or associated with an organization to exploit attack vectors. The threat actor's intent is to harm an organization or profit from access to an organization's resources.

Type	Description
Virus and harmful code attacks	Tools used by threat actors to disrupt company business, compromise data, or hurt the company's reputation
Unethical gathering of competitive information	This is also known as corporate espionage. The goal is to obtain proprietary information to obtain a competitive advantage or steal clients.

Incident Response Process

A cybersecurity incident refers to either a successful or attempted violation of the security properties of an asset, compromising its confidentiality, integrity, or availability. Incident response (IR) policy sets the resources, processes, and guidelines for dealing with cybersecurity incidents. Management of each incident should follow a process lifecycle. CompTIA's incident response lifecycle is a seven-step process:

Preparation — makes the system resilient to attack in the first place. This includes hardening systems, writing policies and procedures, and setting up confidential lines of communication. It also implies creating incident response resources and procedures.

Detection — discovers indicators of threat actor activity. Indicators that an incident may have occurred might be generated from an automated intrusion system. Alternatively, incidents might be manually detected through threat hunting operations or be reported by employees, customers, or law enforcement.

Analysis — determines whether an incident has occurred and performs triage to assess how severe it might be from the data reported as indicators.

Containment — limit the scope and magnitude of the incident. Incident response aims to secure data while limiting the immediate impact on customers and business partners. It is also necessary to notify stakeholders and identify other reporting requirements.

Eradication — removes the cause and restores the affected system to a secure state by applying secure configuration settings and installing patches once the incident is contained.

Recovery — reintegrates the system into the business process it supports with the cause of the incident eradicated. This recovery phase may involve restoring data from backup and security testing. Systems must be monitored closely to detect and prevent any reoccurrence of the attack. The response process may have to iterate through multiple phases of identification, containment, eradication, and recovery to affect a complete resolution.

Lessons learned — analyzes the incident and responses to identify whether procedures or systems could be improved. It is imperative to document the incident. Outputs from this phase feed back into a new preparation phase in the cycle.

Incident response likely requires coordinated action and authorization from several departments or managers, which adds further complexity. The IR process is focused on cybersecurity incidents. There are also significant incidents that pose an existential threat to company-wide operations. These major incidents are handled by disaster recovery processes. However, a cybersecurity incident might lead to a major incident being declared.

9.1.3 Isolate and Contain (Lesson Video)

Transcript:

Today, we'll be discussing isolation, containment, and segmentation within network security. Let's get started.

Isolation is limiting the ability of a compromised asset or application from doing more harm to the network or its assets. This can be accomplished in a few different ways. One way is to practice process isolation. This ensures that if a process is compromised, only the resources used by that process are at risk. This practice applies to operating systems as well as RAM. In other words, it prevents any process that is limited by access bounds from accessing the resources of another process. This is a trait of a stable operating system. Isolation is considered a preventative security measure since it's implemented before an event is detected.
Containment is the first step after an event has been detected and identified. This action can take a few forms. An IT admin may disconnect a machine from the network by simply unplugging the Ethernet cable or disabling the NIC. If this network is connected to other networks, this connection may be terminated. The decision to disconnect must be weighed against the amount of data being compromised and the potential loss of forensic evidence. No matter what, the goal of containment is to limit the damage potential of malicious activity.
Containment requires action. Once an IT security analyst detects and identifies a malicious event, they must act. In this scenario, the analyst is monitoring a physical server that must be manually disconnected from network. This means the on-site IT Admin must jump into action as quickly as possible. Time is of the essence since this event threatens the physical server and also the servers in the branch office. This is because the two networks are connected via a VPN. Containment requires that the damage be limitedâ€”even if it means taking a server down.
Segmentation is a strategic network design. The concept is simple: keep sections of a network separated so that malicious actors can't pivot within a network. Segmentation can be accomplished through VLANs, software-defined networks, switches, subnetting, or even physical segmentation.
But simply being on a different subnet is not enough. Rules must be implemented to control what kind of communications can occur between assets on the network. Many times, a network admin will create a DMZ. This a virtual area where assets are kept separate from internal network assets. A network with a DMZ may have a single firewall or two firewalls depending on how secure this segment needs to be. No matter the topography, access between the DMZ and the internal network is secure and controlled.
That's it for this lesson. We discussed isolation and how it's used to protect a network. Next, we talked about containment, which is the first action taken once an event has been detected. We ended by discussing network segmentation and how it can prevent unauthorized access.

9.1.4 Isolate and Contain Facts

This lesson covers the following topics:

- Isolation, containment, and segmentation
- Security Orchestration, Automation, and Response (SOAR)
- Incident plans

Isolation, Containment, and Segmentation

Data, whether good or malicious, must be handled correctly. You can use isolation and containment for malicious or suspect data. You can use segmentation as a strategic network architecture tool to prevent outside data from accessing internal network appliances.

Strategy	Description
Isolation	Isolation limits the ability of a compromised process or application to do more harm to the network or its assets. One way to protect the network is process isolation. This ensures that if a process is compromised, only the resources used by that process are at risk.

Strategy	Description
Containment	Containment is the first step after an event has been detected and identified. This action can take a few forms. You can disconnect a machine from the network by unplugging the Ethernet cable or disabling the NIC. If a network is connected to other networks, you can terminate those connections.
Segmentation	Segmentation is a strategic network design. The concept is simple: separate the network sections so malicious actors cannot pivot within a network. You can segment using VLANs, software-defined networks, switches, subnetting, or physical segmentation. Being on a different subnet is not enough. You must implement rules to control the kind of communications that occur between assets on the network. You can also create a demilitarized zone (DMZ). It is a virtual area where you separate assets from internal network assets. Depending on how secure the segment needs to be, a network with a DMZ may have a single firewall or two firewalls. No matter the topography, access between the DMZ and the internal network is access-controlled.

Security Orchestration, Automation, and Response (SOAR)

SOAR is a platform to compile security data generated by different security endpoints. This collected information is then sent to a security analyst for further action. SOAR frees an analyst from constantly receiving security alerts as they are generated. Analysts can use parameters to automate solutions for security incidents that meet specific criteria. SOAR:

- Gathers alert data and places it in a specified location.
- Facilitates application data integration.
- Facilitates focused analysis.
- Creates a single security case.
- Allows for multiple playbooks and playbook step automation.

Incident Plans

As part of the incident response process, you can use playbooks and runbooks together to achieve a more effective response that can be automated and include tasks automatically assigned to analysts to complete. These two plans can also help to meet and comply with regulatory frameworks like GDPR or NIST if necessary.

Plan Type	Description
Runbooks	Runbooks are a condition-based series of protocols you can use to establish automated processes for security incident response. Assessment, investigation, and mitigation are accelerated using a runbook. Even though processes are automated, human analysis is still used in some cases.
Playbooks	A playbook is a checklist-style document specifying how to respond to a threat or incident. The steps are listed in the order to be performed. A playbook ensures a consistent approach to security issues.

9.1.5 Summarize Incident Response Procedures

9.1.6 Practice Questions (Section Quiz)

q_incident_resp_analyze_01_secp8

A large organization's cybersecurity incident response team receives an alert indicating potential threat actor activity on one of its network servers.

What should be the team's immediate action based on the incident response life cycle?

Answers:

Wait for more alerts to confirm the incident before taking any action.

Immediately disconnect the affected server from the network to isolate it.

***Analyze the alert and its context to determine whether a genuine incident has occurred.**

Notify the executive decision-maker to authorize actions before proceeding.

Explanation:

When receiving an alert, the first responder's immediate action is to determine whether a genuine incident has occurred. This action involves investigating the data reported and assessing the severity of the situation before taking further action.

Waiting for more alerts to confirm the incident is not advisable. Immediate analysis is necessary to assess the severity of the situation and determine whether it is a genuine incident.

Immediately disconnecting the affected server from the network might disrupt business operations and hinder the investigation. The organization should analyze the incident before taking such drastic action.

Notifying the executive decision-maker to authorize actions before proceeding would delay acting on the threat and potentially cause serious damage to the network server.

q_incident_resp_analyze_02_secp8

The computer incident response team (CIRT) has informed the executives of a large financial institution of unusual network activity, indicating a potential breach.

Which phase of the incident response lifecycle involves investigating the reported unusual network activity to determine whether a genuine security incident has occurred and assessing the severity of the situation?

Answers:

***Analysis**

Preparation

Containment

Eradication

Explanation:

During the incident response life cycle, the analysis phase is where the CIRT must investigate the reported unusual network activity to determine whether a genuine security incident has occurred and to assess the severity of the situation.

The preparation phase focuses on setting up security measures, policies, and communication lines to prevent incidents.

The containment phase occurs after the analysis phase and involves limiting the scope and magnitude of the incident to prevent further damage. It focuses on securing data and minimizing the immediate impact on customers and business partners.

The eradication phase follows containment and involves removing the intrusion tools and unauthorized changes from the systems after the CIRT has contained the incident.

q_incident_resp_containment_01_secp8

A cybersecurity analyst in a multinational corporation is responsible for sensitive customer data and proprietary information and is now dealing with a security breach.

The team is managing the incident response process using the CompTIA incident response life cycle. The team has just completed the third step in the process.

What must the team do next?

Answers:

Preparation

Detection

Analysis

***Containment**

Explanation:

Containment is the fourth step in the CompTIA incident response lifecycle, as it involves limiting the scope and magnitude of the incident, securing data, and notifying stakeholders.

The first step in the CompTIA incident response life cycle is preparation, which focuses on making the system resilient to attacks and setting up incident response resources and procedures.

The second step in the CompTIA incident response life cycle is detection, which discovers indicators of threat actor activity. Indicators that an incident may have occurred might be generated from an automated intrusion system. Alternatively, incidents might be manually detected through threat hunting operations or be reported by employees, customers, or law enforcement.

The third step in the CompTIA incident response life cycle is analysis, which determines whether an incident has occurred and performs triage to assess how severe it might be from the data reported as indicators.

q_incident_resp_containment_02_secp8

A multinational manufacturing company's Chief Information Officer (CIO) heavily depends on its digital infrastructure for managing supply chains and production.

Recently, the IT team detected some unusual activity on the company's network that might indicate a potential malware infection.

The incident response team is now following the CompTIA incident response life cycle to tackle the issue. As the CIO guides them through the different phases, they need to highlight the primary objective of the containment phase.

Which of the following options will they be highlighting?

Answers:

Restoring the affected system to a secure state.

Identifying indicators of threat actor activity.

***Securing data and limiting the immediate impact.**

Analyzing the incident and improving procedures or systems.

Explanation:

The primary aim of the containment phase is to secure data while limiting the immediate impact on customers and business partners. This phase involves preventing the incident from spreading further and notifying stakeholders.

Restoring the affected system to a secure state is the goal of the eradication phase, which comes after containment and involves removing the incident's root cause.

Identifying indicators of threat actor activity is the goal of the detection phase, where the incident response team discovers signs of a cybersecurity incident through automated systems, manual detection methods, or reports.

Analyzing the incident and improving procedures is the goal of the "lessons learned" phase, where organizations learn from the incident and make necessary changes to prevent future incidents.

q_incident_resp_containment_03_secp8

What is the primary goal of the containment phase of cybersecurity incident management during an incident response lifecycle? (Select two.)

Answers:

Remove all traces of the incident from affected systems.

Analyze the incident and responses to identify whether procedures or systems could be improved.

***Limit the immediate impact of the incident while securing data and notifying stakeholders.**

***Notify stakeholders and identify other reporting requirements.**

Reintegrate the system into the business process it supports with the cause of the incident eradicated.

Explanation:

Containment focuses on:

Limiting immediate impact of the incident from spreading further and minimizing its impact on both data and business operations.

The necessity of notifying stakeholders and identifying other reporting requirements.

The eradication phase involves removing the cause of the incident and restoring affected systems to a secure state.

The lessons learned phase involves analyzing the incident and responding to identify whether procedures or systems could be improved.

The recovery phase involves reintegrating the system into the business process it supports with the cause of the incident eradicated.

q_incident_resp_corporate_espionage_secp8

You are a cybersecurity analyst at a tech startup. Recently, you've noticed an unusual pattern of data access requests from a competitor's IP address. The requests are specifically targeting your company's proprietary algorithms and customer databases.

Based on this information, which type of security incident is MOST likely occurring?

Answers:

Employee errors

Unauthorized act by an employee

External intrusion attempts

Virus and harmful code

Corporate espionage

Explanation:

Unethical gathering of competitive information (or corporate espionage) is the correct answer. The scenario describes a deliberate attempt from a competitor to access proprietary information, which is characteristic of corporate espionage.

Employee errors is incorrect because the scenario describes a deliberate attempt to access proprietary information, not an unintentional action by an employee that causes damage or leaves network systems vulnerable to attack.

Unauthorized act by an employee is incorrect because the source of the data access requests is from a competitor's IP address, which suggests that the threat is external, not internal.

External intrusion attempts is incorrect because while the data access requests are coming from an external source, the specific targeting of proprietary algorithms and customer databases suggests a motive beyond a general intrusion attempt.

Virus and harmful code attacks is incorrect because the scenario does not mention any signs of malicious code or software, such as system crashes, slow performance, or unexpected pop-ups, which are typically associated with virus and harmful code attacks.

q_incident_resp_external_intrusion_attempts_secp8

You are a cybersecurity analyst at a large corporation. You receive an alert from your intrusion detection system (IDS) indicating a sudden spike in network traffic and multiple failed login attempts on a server containing sensitive data.

The source IP addresses are from various foreign locations.

Based on this information, which type of security incident is MOST likely occurring?

Answers:

Employee errors

Unauthorized act by an employee

***External intrusion attempts**

Virus and harmful code attacks

Unethical gathering of competitive information

Explanation:

External intrusion attempts is the correct answer. The scenario describes a deliberate attempt from various foreign locations to gain unauthorized access to a server, which is characteristic of an external intrusion attempt.

Employee errors is incorrect because the scenario describes a deliberate attempt to gain unauthorized access to a server, not an unintentional action by an employee that causes damage or leaves network systems vulnerable to attack.

Unauthorized act by an employee is incorrect because the source IP addresses are from various foreign locations, which suggests that the threat is external, not internal.

Virus and harmful code attacks is incorrect because the scenario does not mention any signs of malicious code or software, such as system crashes, slow performance, or unexpected pop-ups, which are typically associated with virus and harmful code attacks.

Unethical gathering of competitive information is incorrect because the scenario does not provide any information suggesting that the goal of the intrusion attempts is to obtain proprietary information for competitive advantage or client theft.

q_incident_resp_incident_secp8

What is the BEST definition of a security incident?

Answers:

Criminal activity

***Violation of a security policy**

Compromise of the CIA

Interruption of productivity

Explanation:

The best definition of a security incident is a violation of a security policy.

Criminal activity, compromise of the CIA, and productivity interruptions are all violations of security policy. They are specific examples of security incidents rather than a universal definition.

q_incident_resp_lessons_learned_secp8

The leader of the cybersecurity team for a major e-commerce company recently encountered a major data breach that led to the exposure of customer payment details. The team has now contained the breach and is moving toward the final phase of the incident response cycle.

What is the team's primary objective in this phase?

Answers:

Identify stakeholders and reporting requirements

Eradicate the cause of the incident

Restore the affected system to a secure state

***Analyze the incident and improve procedures or systems**

Explanation:

The final phase of the incident response lifecycle entails "lessons learned," which allows the organization to learn from the incident and make necessary changes to prevent similar incidents in the future.

During the containment phase in the incident response lifecycle, the incident response team identifies stakeholders and reporting requirements, limiting the scope of the incident and notifying relevant parties.

The eradication phase follows the containment phase. The primary goal is to eradicate the cause of the incident and involves removing the root cause and restoring affected systems.

The recovery phase in the incident response lifecycle aims to restore the affected system to a secure state and reintegrate it after containing and eradicating the incident.

q_incident_resp_process_secp8

An organization's computer incident response team (CIRT) receives an alert that shows possible malicious activity on a critical server within the network, and they initiate the CompTIA incident response process.

The team follows the incident response lifecycle to address the situation, which involves several key steps.

What order must the CIRT follow when performing the CompTIA incident response process?

Answers:

Preparation, analysis, isolation, containment, recovery

Detection, analysis, eradication, restoration, improvement

***Detection, analysis, containment, eradication, recovery**

Isolation, analysis, restoration, eradication, improvement

Explanation:

This sequence of detection, analysis, containment, eradication, and recovery represents the CompTIA incident response lifecycle steps. It ensures the CIRT properly addresses the incident while minimizing damage and potential recurrence.

Isolation is a step that is not part of the CompTIA incident response lifecycle. Additionally, isolation is often a containment technique rather than a standalone step.

q_isolate_contain_books_01_secp8

You would like to enhance your incident-response process and automate as much of it as possible.

Which of the following elements would you need to include? (Select two.)

Answers:

***Runbooks**

***Playbooks**

Whitelisting

Blacklisting

Quarantining

Explanation:

The correct answers are runbooks and playbooks. Runbooks are a condition-based series of protocols you can use to establish automated processes for security-incident response. A playbook is a checklist-style document that specifies the steps to be taken in response to a threat or incident. The steps are listed in the order to be performed. A playbook ensures a consistent approach to security issues.

Whitelisting allows an IT admin to control the applications, IP addresses, URLs, email addresses, etc. that are allowed onto the network.

Blacklisting is the opposite of whitelisting.

Quarantining is when antivirus software finds a malicious item and quarantines it in a special folder.

q_isolate_contain_books_02_secp8

The IT security team of a medium-sized company has identified various attack vectors and threat scenarios that could impact the organization. To enhance their incident response capabilities, the team is developing incident response playbooks.

Additionally, the team is researching external considerations such as legal and regulatory requirements, the potential involvement of law enforcement, and the need to notify affected customers in the event of a data breach.

What is the primary purpose of developing incident response playbooks in the company's cybersecurity strategy?

Answers:

***To provide predefined steps and procedures to respond effectively to cybersecurity incidents.**

To identify potential attack vectors and threat scenarios in the organization.

To ensure compliance with legal and regulatory requirements in case of a security breach.

To prevent cybersecurity incidents from occurring in the company's IT infrastructure.

Explanation:

The primary purpose of developing incident response playbooks is to provide predefined steps and procedures to respond effectively to cybersecurity incidents. Playbooks outline the necessary actions, roles, and responsibilities to minimize the impact of security breaches and handle incidents in a structured manner.

Incident response playbooks involve identifying potential attack vectors and threat scenarios, but the primary focus is establishing a response plan.

Although incident response playbooks address legal and regulatory requirements, the main goal is to outline the response procedures, not solely compliance measures.

Incident response playbooks are for responding to cybersecurity incidents after they occur, not to prevent them proactively. Preventive measures are usually part of a separate cybersecurity strategy.

q_isolate_contain_containment_secp8

You have detected and identified a security event. What is the first step you should complete?

Answers:

***Containment**

Isolation

Segmentation

Playbook

Explanation:

You would choose containment. Containment is the first step to complete after an event has been detected and identified.

Isolation limits the ability of a compromised process or application to do more harm to the network or its assets.

Segmentation is a strategic network design. The concept is simple: keep the sections of a network separated so that malicious actors cannot pivot within a network.

Playbooks are part of an incident-response plan. Playbooks can automate responses.

q_isolate_contain_isolate_secp8

A cyber security analyst at a multinational corporation detects abnormal network activities that indicate a possible security breach. The analyst investigates and confirms that an unauthorized person has accessed sensitive customer information. The incident response team must act quickly to contain the breach and stop further data loss.

What should the initial responder do first?

Answers:

***Disconnect the affected server from the network to isolate it from the production environment.**

Notify law enforcement authorities about the incident for immediate action.

Restore affected systems from secure backups to eliminate the threat.

Initiate a threat hunting exercise to discover evidence of TTPs proactively.

Explanation:

The first responder's primary goal is to contain the incident. Isolating the affected server from the network prevents further communication with the attacker and protects other systems from potential compromise.

Notifying law enforcement might be necessary in certain situations, but the immediate focus should be on containment and minimizing the impact of the incident.

While restoring affected systems from backups is an important step in the recovery phase, it should occur after containment to ensure the team has neutralized the threat.

Threat hunting is a proactive approach to identifying potential threats before they become incidents. In this case, the team has already detected the incident, so immediate containment is the priority.

q_isolate_contain_isolation_secp8

You need to limit a compromised application from causing harm to other assets in your network.

Which strategy should you employ?

Answers:

***Isolation**

Segmentation

Containment

SOAR

Explanation:

The correct answer is isolation. Process isolation is one way to protect a network. It ensures that if a process is compromised, only the resources that are used by that process are at risk.

Segmentation is a strategic network design. The concept is simple-keep the sections of a network separated so that malicious actors cannot pivot within a network.

Containment is not a preemptive strategy. Containment is something you do after an event has occurred.

SOAR is a platform to compile security data generated by different security endpoints. This compiled information is then sent to a security analyst for further action.

q_isolate_contain_segmentation_01_secp8

You need to limit the impact of a security breach for a particular file server with sensitive company data.

Which strategy would you employ?

Answers:

***Segmentation**

Isolation

Containment

SOAR

Explanation:

The correct answer is segmentation. You can segment using VLANs, software-defined networks, switches, subnetting, or even physical segmentation.

Isolation limits the ability of a compromised process or application to do more harm to the network or its assets.

Containment is the first step after an event has been detected and identified. Segmentation is preventative.

SOAR is a platform to compile security data generated by different security endpoints.

q_isolate_contain_segmentation_02_secp8

The security manager of a multinational organization is on a mission to apply security principles to a newly planned regional office that will connect with its existing global infrastructure. This task aims to minimize the attack surface and construct suitable security zones.

While developing the network architecture for the new office, what primary security aspect must the manager prioritize to certify the efficiency of the security zones and reduce the organization's attack surface?

Answers:

Enabling ports

Increasing redundant network paths

Setting up Single Sign-on (SSO)

***Implementing network segmentation**

Explanation:

Network segmentation involves dividing a network into smaller parts, creating distinct security zones. Network segmentation limits an attacker's ability to move laterally within a network, thereby minimizing the attack surface.

Enabling more ports unnecessarily or without adequate security measures increases the potential entry points for malicious actors, thereby enlarging the attack surface rather than minimizing it.

While redundancy increases network availability, it does not directly contribute to reducing the attack surface or defining security zones.

While SSO can enhance user experience and streamline access management, it does not directly minimize the attack surface or contribute to defining security zones.

q_isolate_contain_segmentation_03_secp8

A distribution company bidding on a national contract must demonstrate that its networks are secure and response times are adequate to stop attacks from reaching the client company.

What can support the company in meeting this requirement?

Answers:

***Segmentation**

Isolation

Hardening

Patch management

Explanation:

A segmented network divides systems into separate segments or subnets, each with distinct security controls and access permissions. This type of segmentation significantly complicates an attacker's work, giving an organization more time to detect and respond.

Isolation segregates individual devices within a network to limit their interaction. While isolation can help protect the network, segmentation provides better protection.

Hardening is the process of making an operating system or application more secure. It does not directly affect the network.

Patch management is a process for reviewing, scheduling, or automatically installing patches. Patch management protects devices, not the network.

q_isolate_contain_soar_secp8

As a security analyst, you are looking for a platform to compile all your security data generated by different endpoints.

Which tool would you use?

Answers:

***SOAR**

MDM

MAM

GDPR

Explanation:

The correct answer is SOAR (Security Orchestration, Automation, and Response). This compiled information is sent to a security analyst for further action. SOAR frees an analyst from constantly receiving security alerts as they are generated. Analysts can use parameters to automate solutions for security incidents that meet certain criteria.

An MDM is for managing mobile devices. It is not for all endpoints.

An MAM allows you to manage mobile apps on all sorts of devices, but it does not allow you to compile endpoint data.

GDPR (General Data Protection Regulation) is a framework in the EU for data protection and privacy.

9.2 Log Management

As you study this section, answer the following questions:

- What does a security information and event management (SIEM) system do?
- Why are trends important for network management?
- What part does event correlation play in a SIEM?
- How do IT security teams use alerts?

In this section, you will learn to:

- Save captured files with Wireshark.
- Use Elasticsearch, Logstash, Kibana.
- Use NetworkMiner
- Configure remote logging on Linux
- Log events on pfSense.

Key terms for this section include the following:

Term	Definition

SIEM	A software tool used to compile and examine multiple data points gathered from across a network.
Sensor	A device that gathers data from a device or system. It provides the collected data to a monitoring system.
Trend	Patterns of activity discovered and reported to the SIEM.
Sensitivity	Customized threshold for sensor data that is sent to the SIEM.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Monitoring 3.2 Given a scenario, apply security principles to secure enterprise infrastructure. Infrastructure considerations Intrusion prevention system (IPS)/intrusion detection system (IDS) Sensors 4.1 Given a scenario, apply common security techniques to computing resources. Secure baselines Establish 4.3 Explain various activities associated with vulnerability management. Identification methods Vulnerability scan Analysis Common Vulnerability Enumeration (CVE) 4.4 Explain security alerting and monitoring concepts and tools. Activities

	Log aggregation Alerting Tools Security information and event management (SIEM) NetFlow Vulnerability scanners 4.9 Given a scenario, use data sources to support an investigation. Log data Firewall logs Application logs Endpoint logs OS-specific security logs IPS/IDS logs Network logs Metadata Data sources Dashboards Packet captures
TestOut Security Pro	5.1 Implement Logging and Auditing Implement Logging and Auditing 5.1.2 Enable Device Logs 5.2 Assessment techniques 5.2.1 Implement intrusion detection 5.2.3 Scan for vulnerabilities

9.2.1 Security Information and Event Management (Lesson Video)

Transcript:

In this video, we'll discuss Security Information and Event Management, or SIEM, which is a tool used to compile and examine multiple data points gathered from across a network. We'll also explore log management.

Monitoring a network requires experience and solid tools. One common network security tool is a scanner that can identify vulnerabilities and recommend remediation steps. The scan delivers output information to the IT administrators via the SIEM dashboard. The interval between scans is set by the IT department. This tool also scans servers, firewalls, switches, software programs, and even security cameras and wireless access points.
There are many versions of SIEMs. Each one has different features and benefits, but all SIEMs have several features in common. One of these is the dashboard. These are generally customizable information screens that show real-time security and network information. This allows the IT security team to monitor events as they occur on the network.
Sensors play a vital role in monitoring and securing a network. These sensors are set up at critical endpoints, services, and other vulnerable locations. They're programmed to send customized alerts to the SEIM if certain parameters are reached or exceeded.
Before sensors are deployed, the IT security team sets their sensitivity. The benefit to variable sensitivity settings is the ability to customize the data that's sent to the SIEM. Not every company has the same needs, and that's what makes sensor-sensitivity customization so beneficial.
Trends are patterns of activity discovered and reported to the SEIM. This is how baselines are established. These trends help security analysts decide if reported activity is normal or outside of the baseline. Trends that don't fit previously recorded ones can be investigated. As the IT security team investigates and documents these trends, it becomes easier for them to quickly spot an anomaly that may signal a security concern.
Alerts are the SIEMs' way of letting the IT team know that a pre-established parameter has been contravened. The alert is intended to get the attention of the IT person who's monitoring the network. A best practice in this area is 24-hour monitoring. This means that weekends, holidays, and early hours are all filled. Hackers don't keep normal hours, and network equipment can break at the most inconvenient times!
Event correlation is a critical part of using a SIEM solution. The software gathers data from log files, system applications, network appliances, and other endpoints in order to analyze it. This work is tedious, and people are inefficient at it. That's why the event-correlation feature is valuable. Not only does it gather data, but it analyzes and compares known malicious behavior against the aggregate data so that events aren't missed.
That's it for this lesson. In this lesson, we learned about SIEMs and the features that make them a critical network component. These include customizable sensors and their placement. Trends help IT teams to establish baselines or norms for their network. Alerts provide critical information to the network monitor so that appropriate action can be taken. Event correlation automates a very laborious process and analyzes and compares aggregate data to known security behaviors and events, ensuring that nothing gets by undetected. All this is delivered to the IT team via a customizable dashboard for easy examination.

9.2.2 Log Management (Lesson Video)

Transcript:

In this video, we'll explore log categories, specific logs you should know, and open-source tools that are used to aid IT security teams.

Every network generates dozens and dozens of logs. Network logs tell us what's coming into and leaving our network. Every network appliance and almost every application produces logs which can be used for a variety of reasons. But more often than not, we use logs for network security.
System logs are produced by an operating system. These logs contain all the information that pertains to that OS like updates, errors, failures, and other system occurrences. This includes information for client computers and servers, which gives IT admins a way to investigate events on individual machines that may interconnect with other machines on the network.
Most applications produce some type of event logging. These logs show application access, crashes, updates, and any other relevant information which could be valuable in determining root-cause analysis. The application may be crashing or not performing correctly, and this could be tied to suspicious activity that may indicate malicious intent.
There are several logs that would fall under the security category. There are application-security logs, event-security logs, and security logs for specialty applications like IDS/IPS, endpoints, firewalls, routers, and switches. Also, logs for security cameras and wireless or physical access points are included under the security log umbrella.
There's no doubt that web server logs are one of the most tedious of all logs to parse. But web servers can be prime targets for hackers, so it's important to know who's interacting with your server and what they're attempting to do. Most web engines like IIS, Tomcat, Web Sphere, and NGINX have some level of server logging. These logs can tell you

exactly when users log onto your site and what their location is. They also give you some information on attempted attacks.
Domain Name System, or DNS, has been around for a long time. When DNS was designed, network security wasn't a priority. Over time, malicious actors started using DNS-targeted attacks. These attacks have the potential to be disruptive and quite expensive. With DNS logging, you can track updates and choose to approve or deny them. DNS also produces query logs that detail which requests are being handled by which instance. Rate limiting is another valuable tool that limits response rate. Analyze these logs to see when rate limiting was used and for what purpose. Client IP, record requests, flags, and other metadata can be included in these logs.
Dump files are created when an application, OS, or other computer function stops abruptly. The information that's stored in memory at the time is dropped into a file for later analysis. These files help IT admins perform root-cause analysis because they provide the state that the application was in when it crashed, error codes, and other clues as to what happened previous to the application failure. They can also give clues as to the crash's origin. This could be something as commonplace as a bad driver or hardware component, or it may prove, unfortunately, to be the result of a malicious act.
Authentication logs are vital to a network's security. Authentication servers may be Active Directory-based or OpenLDAP depending on your network structure. It's critical to know who may be poking around your network, so token requests, authentication failures, or failed logins on expired accounts are all stored on these authentication logs for you to view.
Voice over Internet Protocol, or VoIP, has become a common network application. With a high implementation rate comes attention from malicious actors. As with any network application, there are vulnerabilities that can be leveraged, so in order to defend it there needs to be a way to access information about what's happening at any given time.
Session Initial Protocol, or SIP, is the standard in VoIP calls. The key to tracking attacks against a VoIP system is understanding SIP and being able to parse its logs. These logs contain key information about where a call was initiated and what the communication's intent was. These facts help the IT security team create a stronger SIP security posture.
Syslog is short for System Logging Protocol. This protocol sends system logs and event messages to a server designated by the system administrator. It collects logs from various appliances and sends them to the syslog server where they can be reviewed and analyzed.
Rsyslog is an open-source tool created for use in Linux networks. It stands for rocket-fast system-log processing. It gets its name because of its ability to send a million messages per second to a local server. This tool's benefit is its diversity. It's capable of multi-threading by leveraging multiple security protocols like TCP, TLS, and others. It also allows for output-format customization.
Syslog-ng is a robust log-aggregating software for multiple platforms, including Windows. This tool increases the quality of the log data that's sent to your SIEM. It also facilitates lightning-fast log searches by using full-text queries and collects logs without the installation of server agents.
Journactl is a Linux tool that gathers the logs produced by systemd, a system that's the basis for many Linux components. This command is used in Bash to parse logs that've been collected by systemd. The results are presented in the syslog format and are ordered oldest to newest by default. This can be changed by using the â€“r flag. Each line shows the date, server hostname, process name, and any messages. If you're more comfortable using a command line tool, this is for you. Because it's a CLI tool, there are many key commands at your disposal to get you your logs quickly.
Nxlog is an open-source log-collector application. It uses log-collector agents to gather and send log data to a log server, which can itself be set up using nxlog. This application is available for both Windows and Linux. The Community edition supports multiple SIEM applications and works with Windows Event Viewer and syslog.
That's it for this lesson. In this lesson, we learned about different log categories like network logs, security logs, and application logs. We also learned about specific logs and how you can use them for better security. Finally, we learned about open-source tools that assist you in collecting and organizing log data and metadata. Understanding logs is truly the key to better security.

9.2.3 SIEM and Log Management Facts

This lesson covers the following topics:

- Security information and event management

Security Information and Event Management

A security information and event management (SIEM) system combines security information management (SIM) and security event management (SEM) functions into one security management system.

Security information and event management tools compile and examine multiple data points gathered across a network. The following table describes SIEM components.

Component	Description
Vulnerability scan output	Monitoring a network requires experience and solid tools. One tool standard to network security is a scanner that can identify vulnerabilities and recommend remediation steps. This tool scans servers, firewalls, switches, software programs, security cameras, and wireless access points. The scan delivers the output to IT admins via the SIEM dashboard. The interval between scans is set by the IT department.
SIEM dashboards	The dashboard is a common component of all SIEM systems. The dashboard consists of customizable screens showing real-time security and network information. The information in real-time allows the IT security team to monitor and respond to events on the network effectively .
Sensors	Sensors are a vital part of monitoring and securing a network. Sensors are set up at critical endpoints, services, and other vulnerable locations. These sensors are programmed to send customized alerts to the SEIM if specific parameters are not within the acceptable range.
Sensitivity	The IT security team sets the sensitivity level when the sensors are deployed. The benefit of variable sensitivity settings is the ability to customize the data sent to the SIEM. Not every organization will have the same needs in network monitoring.
Trends	Trends are patterns of activity discovered and reported to the SIEM. This is how baselines are established. Trends help security analysts decide whether reported activity is normal or outside the baseline. The security group can investigate trends that do not fit previously recorded information. As the IT security team investigates and documents these trends, it becomes easier for the team to spot a trend that may signal a security event quickly.
Alerts	Alerts are the SIEM's way of informing the IT team that a pre-established parameter is not within the acceptable range. The alert is intended to get the attention of the IT person or persons monitoring the network. A best practice in this area is 24-hour monitoring.
Correlation	Event correlation is a critical SIEM component. The software gathers data from log files, system applications, network appliances, etc., and analyzes it. This work is tedious; people are inefficient at it. That's why the event correlation feature is valuable. Not only does it gather the data, but it analyzes and compares known malicious behavior against the aggregated data, increasing the chances of the discovery of security events.

9.2.4 Monitoring Data and Metadata (Lesson Video)

Transcript:

In this lesson, we'll cover the use of bandwidth monitors and metadata from emails, mobile devices, web traffic, and files. We'll also talk about using netflow/sflow echo and Ipfix. Let's get started.

Bandwidth monitors are a type of application that help network admins understand bandwidth usage. The first order of business with these applications is to establish a baseline. The longer a monitor runs, the more data points are created.

When a substantial number of data points are created, a normal usage becomes apparent. Normal bandwidth usage is relative and varies by hours in the day, days of the week, and even weeks within the year. A baseline provides something concrete to compare against current usage or even suspect bandwidth usage.
There are many bandwidth monitoring applications available. There are cloud-based apps, on-premise apps, open source, and paid. Each one looks and functions differently, but each has the same goal, which is to allow easy access to bandwidth monitoring. The key is to learn how to establish a monitoring schedule and how to use the dashboard.
Here, we have some screenshots of different pages within a dashboard. As you can see, this graphic is showing bandwidth usage be the hour and below usage per day aggregated over time. This is the baseline that suspect usage is compared against. The second graphic shows the last hour's usage and the usage for a user-specified time frame. This allows for deeper examination of suspect bandwidth usage.
Email is a great tool that almost everyone uses to communicate. It's also the avenue for the majority of malicious network breaches. Fortunately, email provides metadata so it can be traced. All emails come with a header that contains information about both the sender and recipient. Parts of the headers can be spoofed to give investigators false information. The good news is that there are security devices that put X-headers throughout an emails' headers. These provide the originating email account and IP address, not the spoofed one.
Communications sent via mobile devices are common today. They come from tablets, laptops, smart phones, smart watches, and any other portable device that connects to the internet. These devices send emails and text messages and use apps to allow data and photo sharing. All this data produces metadata that can be used to identify people, places, times, and even deleted data. Pictures can be time stamped and geolocation stamped. Much of this metadata also reveals the origination of the data and the sender.
Websites produce many types of metadata. In fact, the metadata on a user's machine versus the server can be very different. The data on the opposite side of the transmission can help fill in gaps and corroborate findings on the opposing machine. Metadata includes IP addresses, user requests, user downloads, time spent on the site, and even attempts to gain unauthorized access. Web metadata also includes cookies, browser history, and even cached pages. Many times, malicious actors will attempt to obfuscate their real metadata. But the good news is that there are ways of finding the real metadata, especially for trained forensic investigators.
Files produce many types of metadata. The first kind is File Creation (date/time). The File Creation data is the first time the file was written to the storage media it's currently on. This means that the file can be created in a different place and then moved or copied to a new location.
Last written refers to the last time a file was saved for any reason. It could prove that the file was copied from a different location or that the file was altered. Last access is any time a file is touched. When combined with file creation metadata on a user's machine, it can establish the probability that the file was copied from one machine to another. It also introduces the idea that a third device is involved in data copy process. Once the third device is found, it makes obtaining answers much easier.
Network admins are always looking for a way to look at what's happening inside their network. Netflow is a session sampling protocol. This protocol works at Layers 2 through 4. It can exam each data flow that comes through or be set to sample sessions at certain intervals. If you want to sample packets in a broader layer range, then sFlow is your answer. This protocol works on Layers 2 through 7. Unlike Netflow, sFlow examines packets and can only be used in sampling mode. This is stateless packet sampling that provides information efficiently.
IPFIX integrates data that normally goes to Syslog or SNMP information directly in the IPFIX packet, eliminating thee additional services collecting data from each network device. IPFIX has provisions for fields that are variable-length, meaning no ID number restrictions. It came about because of the need for a standardized protocol for internal protocol flows. This data comes from routers, servers, and other network appliances that are mediation systems. The data is formatted and sent to an exporter and then on to a collector.
IPFIX, like NetFlow, looks at flow and the number of packets being sent and received during a given session.
That's it for this lesson. In this lesson, we talked about the different ways metadata is produced. We covered some of the pitfalls of metadata as well as some of the ways to overcome them. And we discussed creating a baseline using network bandwidth monitors. These monitors can alert us to abnormal activity on our network.

9.2.5 Saving Captured Files with Wireshark (Demo Video)

Transcript:

In this demo, we'll show you how to capture network packets. There are times when you might want to capture packets so you can analyze them later. Captured packets can be used by an analyst to profile an application's network traffic or to examine a protocol in more detail.

The two most popular tools to capture packets are Wireshark and TCPDump. Both Wireshark and TCPDump can be used with a variety of operating systems but for this demo we will use Security Onion.
Security Onion is usually set up with a monitor port that captures all packets that it sees. The packets are typically used with tools like Zeek and Snort. This interface can also be used by an analyst for ad-hoc captures.
There are two capture tools in Security Onion for ad-hoc captures. First, we will look at Wireshark. Wireshark is a graphical tool that allows packet capture, but is also an analysis tool. Because Wireshark requires root permissions to capture the packets, we must run Wireshark with elevated permissions.
To do this, go to Applications > Utilities > Terminal. At the command prompt type â€"Ëœsudo dpkg-reconfigure wireshark-common'. You are asked you if you want to allow non-superusers to be able to capture packets. Choose Yes.
Next, run â€"Ëœsudo usermod -a -G wireshark administrator'. This gives the administrative user rights to run Wireshark by adding it to the Wireshark group.
If your account on Security Onion has a different name, use that instead of administrator in the command. Once done, go to the power icon. Then, go down to the other power icon. Choose to restart.
Now the system has rebooted. After logging back into Security Onion, we can start to use Wireshark. We open it by going to Applications > Internet > Wireshark.
To set up a capture, select the interface on Security Onion that is set as the monitor port. In this instance, it is interface enp0s8. After selecting the interface, you can restart the capture by clicking the Shark Fin in the top left of the menu bar. By default, the captured packets will scroll by on the screen.
When we are done capturing, we can press the red stop button. At this point, we can analyze the captured packets as we would any other capture that has been given to us. We can then save that file as a PCAP file for later analysis. We're not going to save it here, but you can see we have the option.
Another way that we can capture packets with Security Onion is to use the command-line tool tcpdump. Let's open a new Terminal session and enter the command â€"Ëœsudo tcpdump -D' to list the possible interfaces that we can capture on. Note that enp0s8 is number 10 on the list.
To capture, run the command â€"Ëœsudo tcpdump -i 10 -w testout.pcap'. This will capture the packets on enp0s8 and write them to a file called testout.pcap. Press CTRL-C to stop the capture.
Now let's print the file to the screen using â€"Ëœtcpdump -r testout.pcap'. Here you can see the output of the file.
We can open this PCAP file in a tool like Wireshark or NetworkMiner. Let's open it in Wireshark. We type, â€"Ëœwireshark testout.pcap' and press Enter. After a second or two the PCAP file is loaded into Wireshark and we can use it to further analyze the packet capture.
That's it for this demo. In this demo we used Wireshark and TCPdump to capture packets and saved them to a PCAP file.

9.2.6 Use Elasticsearch Logstash Kibana (Demo Video)

Transcript:

In this demo, we use the Elasticsearch Logstash Kibana (ELK) stack to store and search security logs created by other tools in Security Onion such as Zeek. We'll use Kibana to review sample logs.

Security Onion is a free and open source Linux distribution for threat hunting, enterprise security monitoring, and log management. It includes Elasticsearch, Logstash, Kibana, Snort, Suricata, Zeek (formerly known as Bro), Wazuh, Sguil, Squert, CyberChef, NetworkMiner, and many other security tools.
First, we click the Kibana link on the Security Onion desktop. By default, Security Onion uses a self-signed TLS certificate. We tell Chromium to allow us to proceed by clicking Advanced and then proceed. On a production system, it would be advisable to install a valid TLS certificate and use the fully qualified name instead of localhost.
The username and password for Security Onion are set during the initial setup. In this case, we set up a user in the Security Operations Center (SOC). We'll login as that user.
After logging in, Kibana defaults to the dashboard page. This could be used as a dashboard for a SOC or as the starting point for a threat hunter. In this case, we want to look at the Network Intrusion Detection System (NIDS) logs, so we click NIDS to drill down.

Here we notice a classification of Attempted Administrator Privilege Gain that we want to look at further. To do that, we hover on that line and click the magnifying glass with the plus sign. This adds it as a filter. To keep the filter as we pivot to other parts of Kibana, we click Actions and pin the filter.
Next, we click Discover to see the logs that match the filter. Notice how the filter has stayed with us because we pinned it. Now we can see each log entry that matched our filter. Notice that the timestamp and other information about the alert is listed.
You can view more information for an event by clicking the arrow to the left of the time stamp. From the description, we can also search for more information about the signature listed. We do this by highlighting it, right-clicking, and choosing to search for it using Google.
Some signatures are based on Common Vulnerabilities and Exposures (CVEs). CVEs describe potential security issues with certain software or hardware. They provide a common language to evaluate the risk posed. By typing CVE into the search field, you can focus the results on signatures that are tied to a CVE. Researching the CVE can help you determine whether the attempt against the asset could be successful.
Let's look at one more example. But first, let's clear our pinned criteria by going to Actions and clicking Remove. We also need to remove the search term CVE. Once this is done, we'll go back to the dashboard.
Under Zeek Hunting, click HTTP to list alerts relating only to the HTTP protocol. Now we can look at the events based on the HTTP status messages and the methods used. Using the same technique with the magnifying glass as before, let's drill into the Forbidden messages.
Again, we want to pin the filter so that it stays with us as we move through Kibana. Notice that this time we are using a different way to pin the filter. This method allows you to pin or unpin individual filters instead of all filters.
In the details of one of the events, we can see the Uniform Resource Identifier (URI) that the attacker tried to access. We can also see the user agent of the attacker. In this case, our attacker didn't mask the fact that a common tool for finding website vulnerabilities called Nikto was being used.
That's it for this demo. In this demo we went over a few of the ways the power of Elasticsearch can be harnessed to help the security professional.

9.2.7 Use NetworkMiner (Demo Video)

Transcript:

NetworkMiner is a Network Forensic Analysis Tool. It is able to take a PCAP file and analyze it for clues about the hosts and protocols on the network at the time of the capture.

To open NetworkMiner in Security Onion, go to Applications, Other and then NetworkMiner. Once it opens, go to file open and then choose the pcap file that you want to analyze. NetworkMiner will then automatically analyze the file and present the information in a series of tabs.
The first tab is the hosts tab. This tab will list all of the unique hosts by IP addresses that were found in the PCAP file. Each host can be opened by clicking on the plus sign next to it revealing the information about the host that NetworkMiner knows. For example, if I open 172.28.24.3 I will get it's IP, MAC Address, information about the number of packets it sent and received, and any details about the host. In this case, if I click on the Host Details, I will see the User-Agent string of the browser the host used.
In the next tab, you will find the files that were transferred between hosts during the packet capture. Right-clicking on a file will allow you to open the file or calculate the file's hash. For this capture, I created a fake virus using the EICAR test string. This string when placed in a file will be detected as a virus by most vendors. I can choose to open the folder where the virus.exe is stored and then open it using gedit to show the EICAR string.
I can also use the file hash feature to compare the file to known bad file hashes using tools like VirusTotal. First right-click on the file and choose Calculate MD5/SHA1/SHA256 hash. Next highlight the hash you want to compare and press CTRL+C to copy it to the clipboard. In a web browser, open VirusTotal.com and click on the Search link. Then paste the hash into the search box and press enter. This will then compare the hash against known malicious files. In this case, it will be an EICAR test file that most antivirus manufacturers will mark as a virus.
Another tab of interest is the credentials tab. Here any clear text credentials like telnet, FTP, or HTTP will be shown. In this capture, there was an FTP session that we have the credentials for.
The sessions tab shows the unique conversations between hosts.
Any DNS queries and responses that are captured will be decoded in the DNS tab.
The parameters tab shows many potential key-value pairs that may have important information.

Finally, NetworkMiner provides the ability to search a PCAP file for certain keywords. You can do this by clicking on Keywords. Next add a keyword and follow the instruction to reload the case file. The display will now show all of the packets that include the keyword.
That's it for this demo. In this demo we used NetworMiner to view the contents of a captured PCAP file.

9.2.8 Configuring Remote Logging on Linux (Demo Video)

Transcript:

There are many aspects of the syslog daemon that you can modify to customize how your log files are configured. You could have warning messages going into one file, error messages going into another, or separate them based on program. But one of the cool things I think the syslog daemon can do is log to a remote host. This basically allows us to write our logs not only to our local system but also to a log host somewhere else. This is a great benefit for the administrator. You can set up a log host in your network somewhere and all the logs from all the systems you support go into that log host. If somebody's having a problem, instead of having to use SSH in their system to look at their log files, you can have one central location to view them. This can be very helpful in keeping history and tracking your log files. Syslog servers are also very helpful when there has been an attack on your network.

That's what we're going to do here. We're going to configure the syslog daemon to log to a remote host. I have two different Linux systems running. I have a RedHat system running here that will serve as my log host and I have a CentOS system running here that will serve as my log client. The log messages from the CentOS system are going to be saved on our syslog server which happens to be our Redhat system. Keep in mind that different Linux distributions may be slightly different. But the same concepts apply.
Let's configure the log host first. The first thing we need to do is check to see if the rsyslog daemon is running. Type â€˜systemctl status rsyslog'. As you can see everything appear to be in order. We need to change to our root user account, so I'll do a 'su- root' command and enter the password. We're going to edit the configuration file with the â€˜vi /etc/rsyslog.conf' command. As we scroll down, we're going to uncomment some modules. These are the imudp module and the imtcp module. These modules enable the syslog daemon to listen for incoming syslog messages.
Down at the bottom in the rules area we have a template we're going to use for our remote logging. This template isn't here by default so it has to be added. This allows the remote logs to be placed in their own folder by host name and program name. If you don't set up a template, all syslog messages from remote hosts will go in the /var/log/messages of the syslog server. We're going to uncomment this by removing the pound symbols. Let's save it by typing â€˜:wq!'.
For these changes to take effect, we must restart the rsyslog daemon by typing â€˜systemctl restart rsyslog'. Since there will be incoming traffic, we need to modify the firewall to accept incoming messages on port 514.
To do so, type â€˜firewall-cmd --permanent --add-port=514/udp'. When you press Enter you see that it was a success. Now we do the same thing with the TCP protocol with "arrow up" and remove UDP for TCP. The changes won't be active until we reload the firewall with the â€˜firewall-cmd --reload' command.
I'm going to do a netstat to make sure the syslog daemon is listening on port 514.
Type â€˜netstat -tulnp | grep "rsyslo"' and press Enter. This shows us that rsyslogd has port 514 open for UPD and TCP.
We're going to venture over to our CentOS client server that's going to be sending syslog messages to our RedHat server. First, we check to make sure the rsyslog daemon is running by typing "systemctl status rsyslog". We're good to go.
Just like our syslog server, we're going to modify the configuration file for rsyslog. Let's type â€˜sudo vi /etc/rsyslog.conf'. Enter the password and press Enter. We're going to scroll all the way to the bottom to this configuration file.
Just so that I know what this is, I'm going to put a comment in there which is "#syslog server" and type â€˜*.* @@192.168.0.55:514'. The *.* is a wildcard that sends all syslog messages to the syslog server. The IP address listed is the IP of the syslog server and 514 is the port specified to use. We save this with â€˜:wq!' and Enter.
â€˜sudo systemctl restart rsyslog' allows the rsyslog daemon to grab the new settings we just edited.
Now this is the fun part. We will actually get to see a message go over to our syslog server from our CentOS client server. To write a test message, we type â€˜logger this is a test message' and press Enter. Let's use â€˜sudo tail -f /var/log/messages' to see the message locally. Now that we see the local message, let's go over to our syslog server with â€˜cd /var/log/centos-server1/'. I'm going to do a quick â€˜ll' to list what log files have already been written. Since we did a logger command under the TestOut user, we're going to tail that log file to look for our newly created test message with â€˜tail -f testout.log'. As you can see, it's the exact same message that was on our CentOS client server.

That's it for this demo. In this demonstration, we talked about how to configure a log host with Linux on a network. We first looked at the steps you need to complete in order to configure the system that's going to function as the log host to receive logging messages from another system. Then we looked at the log client and configured it to send its log messages not only to its own log files, but also to send a copy to our log host.

9.2.9 Logging Events on pfSense (Demo Video)

Transcript:

In this demo, we're going to spend a few minutes viewing log files on our pfSense security appliance. A log should act as a red flag that something is happeningâ€”potentially something bad. By reviewing your logs on a regular basis, you'll get an idea of the normal traffic on your device. In reality, no one likes to spend hours a day viewing log files. Typically, you would want to configure a syslog server so that all your logs from all devices go to one place to be consolidated for easier analysis.

To view logs on pfSense, we first need to go to Status and then down to System logs. Once we're in System logs, we see the General tab. Like I said, viewing logs isn't the most exciting thing to do, but it's necessary. Under the General tab we can see that we have a time stamp, a process, a PID or process ID, and a message about the log. Let's move on to Gateways.
Under Gateways, you can see that I only have one Gateway on this test system. Looks like my system is grabbing an IP for the WAN network interface. You might be thinking that if it's grabbing an IP from a WAN, shouldn't this be a public IP? That answer is yes, but I do have a test network set up and that network is connected to my regular LAN.
We also have our Routing logs here. Next to that I have my DNS Resolver logs. Finally, we have our Wireless logs. I don't have any Wi-Fi currently configured with this device, so there are no logs for our Wi-Fi.
I'll move over to my Firewall logs. This is generally where you might look for malicious attacks directed toward your network. You'll see information down here with more details. As I scroll down, you can see the different source IPs that triggered the log. This one here, 172.16.1.100, is from my DMZ trying to get out to the WAN.
Dynamic View shows us a bit less detail. Down here I have some WAN traffic on port 5353. As you first get a system set up, you might want to familiarize yourself with the different ports that your firewall is logging. Some ports might be perfectly normal while others might not be. I wasn't familiar with port 5353, but a quick web search told me it's for multicast DNS and is safe. So now I know what it is.
Summary View gives us a bunch of graphs that can be helpful to get a quick visual of things. Here I have my different interfaces. I have three in this deviceâ€”one for my LAN, one for the DMZ, and one connected for the WAN.
I have my protocols and it looks like most of my traffic is UDP.
Down here a little further, I can see the source IPs of what's been triggering the logs. All the 192.168 addresses are from the WAN interface and the 172.16 address is my DMZ.
I have the destination IPs in this next graph. I have my source ports next. Finally, I have destination ports. Here you can see UDP/53 listed. That is my DNS traffic.
So, it's good to get familiar with what your typical traffic looks like. This is called creating a baseline. This isn't covered in this demo, but it should be one of the first things you do when setting up new devices.
My next set of logs are DHCP logs. Here you can see who's getting IPs from the DHCP server. Not only can you see who's getting IPs, but you can see the DHCP acknowledgement, the DHCP requests, and down here you can see that DHCP renewed an IP for one of the clients. All of this is useful information if you're troubleshooting DHCP or need to see what devices are connecting to your network.
I had this device set up as a captive portal at one time so that we could see all the events that are related to it. As a quick review, a captive portal is a web page you're taken to, such as in a hotel or other public place, before you're given access to the internet. You typically must agree to the terms and conditions before being allowed to proceed. You might have to enter a password as well.
The next four tabsâ€”IPSec, PPP, VPN, and Load Balancerâ€”have no log files because I don't have any of those services running on this device. But I do have OpenVPN configured. Down here you can see all of those log files. I do have some NTP logs, or Network Time Protocol logs, here.
The last thing I want to look at is log settings. If you noticed as we were viewing logs, there were only fifty shown. Here is where we can change that if needs be. As I scroll down, you can see other settings that you can configure. You can even reset your logs.

Here at the bottom is where you can configure pfSense to send these logs to a syslog server. I'll check this box to enable it. When I do, I'm presented with some more settings specific to remote logging.
I can be specific about which interface I want to log. I can log all of them, or just one. I could just log my WAN events if that's all I'm concerned about.
I can choose to only have IPv4 logs or only have IPv6 logs sent. Here I would put the IP and port of the remote syslog server that's configured to receive those logs. I don't have a server set up for this demo, but if I did, this is where I would tell pfSense to send the logs to. The format would be something like 10.10.10.100, and it would be port 514. Port 514 is the port that pfSense uses by default to send logs to the syslog server.
Now I would need to tell pfSense what I want to send. I don't like information overload. I only want to send the logs that I actually need. So I might only want firewall events, DNS events, DHCP events, and VPN. In my case, that's OpenVPN. I would then click on Save and, if my syslog server is configured, it'll start to receive log files from my pfSense security appliance.
That's it for this demo. In this demo, we examined the logs on our pfSense security appliance. We viewed many of the different logs that can be collected. We then looked at settings and some things that we can configure. We ended by explaining how to send our logs to a syslog server.

9.2.10 Monitoring Data and Metadata Facts

This lesson covers the following topics:

- Log data
- Metadata
- Data analyzers

Log Data

Log data is a critical resource for investigating security incidents. As well as the log format, you must also consider the range of sources for log files and know how to determine the type of log file that best supports any given investigation scenario.

Event data is generated by processes running on network appliances and general computing hosts. The process typically writes its event data to a specific log file or database. Each event is comprised of message data and metadata:

- Event message data is the specific notification or alert the process raises, such as "Login failure" or "Firewall rule dropped traffic."
- Event metadata is the source and time of the event. The source might include a host or network address, a process name, and categorization/priority fields.

Accurate logging requires synchronization of each host to the same date, time value, and format. Ideally, each host should also be configured to use the same time zone or a "neutral" zone, such as universal coordinated time (UTC).

Windows hosts and applications can use Event Viewer format logging. Each event has a header reporting the source, level, user, timestamp, category, keywords, and hostname.

Syslog provides an open format, protocol, and server software for logging event messages. It is used by a vast range of host types. For example, syslog messages can be generated by switches, routers, firewalls, and UNIX or Linux servers and workstations.

A syslog message comprises a PRI code, a header, and a message part:

- The PRI code is calculated from the facility and severity level.

The header contains a timestamp, hostname, app name, process ID, and message ID fields.

The message part contains a tag showing the source process plus content. The format of the content is application-dependent. It might use space- or comma-delimited fields or name/value pairs.

Metadata

File metadata is stored as attributes. The file system tracks when a file was created, accessed, and modified. A file might be assigned a security attribute, marking it as read-only or a hidden or system file. The ACL attached to a file showing its permissions represents another attribute. Finally, the file may have extended attributes recording an author, copyright information, or tags for indexing/searching.

As metadata is uploaded to social media sites, they can reveal more information than the uploader intended. Metadata can include current location and time, which is added to media such as photos and videos.

Metadata is produced by almost all network activity. Server requests, applications, and email are examples of where metadata can be found. In the context of bandwidth monitors, metadata is used to investigate security-related concerns or incidents. The following table describes three types of metadata.

Type	Description
Email metadata	Email provides metadata that is used to trace email. All emails come with a header containing information about the sender and recipient. Parts of the headers can be spoofed, giving investigators false information. However, there are security devices that put X-headers throughout an email's header. These provide the originating email account and IP address, not the spoofed one.
Mobile metadata	Tablets, laptops, smartphones, smartwatches, and any other device that connects to the internet and can be moved around produces mobile metadata. These devices send emails, text messages, and use apps. All these produce metadata that can be used to identify people, places, times, and even deleted data. Pictures can be timestamped and geolocation stamped. Much of this metadata also reveals the origination of the data and the sender.
Web metadata	Websites produce many types of metadata. The metadata on a user's machine versus the server can differ greatly. The data on both sides of the transmission can help fill in gaps and corroborate findings. Metadata includes IP addresses, user requests, downloads, time spent on the site, and attempts to gain unauthorized access. Web metadata includes cookies, browser history, and cached pages. Many times, malicious actors will attempt to obfuscate their metadata. However, there are ways of finding the actual metadata, especially for trained forensic investigators.

Data Analyzers

Network admins should always look for a way to examine what is happening inside the network. There are several tools to help sift through the tremendous amounts of data generated by network activity. The following table describes some of these tools.

Tool	Description
NetFlow	NetFlow is a feature on Cisco routers. It works at layers 2 – 4. It can examine each data flow that comes through the network or be set to sample sessions at specific intervals.

Tool	Description
sFlow	sFlow is a packet sampling technology that works on layers 2 – 7. Unlike NetFlow, sFlow can only be used in sampling mode. This is a stateless packet sampling that provides information on various layers and does it quickly and efficiently.
IPfix	IPfix directly integrates data that usually goes to Syslog or SNMP. This eliminates additional services collecting data from each network device. IPfix has provisions for variable length fields, meaning there are no ID number restrictions. IPfix addresses the need for a standardized protocol for internal protocol flows. This data comes from routers, servers, and other network appliances that are mediation systems. The data is formatted, sent to an exporter, and then sent to a collector. IPfix, like NetFlow, looks at flow and the number of packets sent and received during a given session.

9.2.11 Practice Questions (Section Quiz)

q_siem_logmgmt_alert_secp8

Which of the following components are the SIEM's way of letting the IT team know that a pre-established parameter is not within the acceptable range?

Answers:

Sensors

Dashboard

Trends

***Alerts**

Explanation:

Alerts are the SIEM's way of letting the IT team know that a pre-established parameter is not within the acceptable range. An alert is intended to get the attention of the IT person monitoring the network. A best practice in this area is 24-hour monitoring.

Sensors are set up at critical endpoints, services, and other vulnerable locations. These sensors are programmed to send customized alerts to the SIEM if certain parameters are not within the acceptable range.

The dashboard consists of customizable information screens that show real-time security and network information.

Trends are patterns of activity discovered and reported to the SIEM.

q_siem_logmgmt_correlation_secp8

As a cybersecurity analyst, you are tasked with identifying a critical component of a Security Information and Event Management (SIEM) system that can analyze and compare known malicious behavior against aggregated data from log files, system applications, and network appliances.

Which component of the SIEM system would be MOST effective for this task?

Answers:

SIEM Dashboards

Alerts

Trends

***Correlation**

Explanation:

Correlation is the correct answer. Event correlation is a critical SIEM component. The software gathers data from log files, system applications, and network appliances and analyzes it. This work is tedious; people are inefficient at it. That's why the event correlation feature is valuable. Not only does it gather the data, but it analyzes and compares known malicious behavior against the aggregated data, increasing the chances of the discovery of security events.

While SIEM dashboards provide a visual representation of real-time security and network information, they do not analyze and compare known malicious behavior against aggregated data. They merely display the information gathered by other components of the SIEM system.

Alerts are the SIEM's way of informing the IT team that a pre-established parameter is not within the acceptable range. However, alerts themselves do not analyze and compare known malicious behavior against aggregated data. They are used to notify the team of potential issues.

Trends are patterns of activity discovered and reported to the SIEM. They help security analysts decide whether reported activity is normal or outside the baseline. However, trends themselves do not analyze and compare known malicious behavior against aggregated data. They are used to analyze the data collected by other components of the SIEM system.

q_siem_logmgmt_dashboard_secp8

As a cybersecurity analyst, you are tasked with implementing a Security Information and Event Management (SIEM) system that allows the IT security team to effectively monitor and respond to events on the network in real-time.

Which component of the SIEM system would be MOST critical for this task?

Answers:

Vulnerability scan output

Sensors

Trends

***SIEM dashboards**

Explanation:

The correct answer is SIEM dashboards. The dashboard is a common component of all SIEM systems. The dashboard consists of customizable screens showing real-time security and network information. The information in real-time allows the IT security team to effectively monitor and respond to events on the network.

While vulnerability scans are essential tools for identifying potential weaknesses in the network, they do not provide real-time monitoring or response capabilities. They provide a snapshot of the vulnerabilities present at a specific point in time, but do not provide continuous monitoring.

Sensors are a vital part of monitoring and securing a network. They are set up at critical endpoints, services, and other vulnerable locations and are programmed to send customized alerts to the SIEM system if specific parameters are not within the acceptable range. However, sensors themselves do not provide a real-time overview of the network events.

Trends are patterns of activity discovered and reported to the SIEM. They help security analysts decide whether reported activity is normal or outside the baseline. However, trends themselves do not provide real-time monitoring or response capabilities. They are used to analyze the data collected over a period of time.

q_siem_logmgmt_sensors_secp8

As a cybersecurity analyst, you are tasked with improving the security posture of your organization. You are considering the implementation of a Security Information and Event Management (SIEM) system.

Which component of the SIEM system would be MOST critical for monitoring and securing network endpoints, services, and other vulnerable locations?

Answers:

Vulnerability scan output

SIEM dashboards

***Sensors**

Trends

Explanation:

Sensors is the correct answer. Sensors are a vital part of monitoring and securing a network. They are set up at critical endpoints, services, and other vulnerable locations and are programmed to send customized alerts to the SIEM system if specific parameters are not within the acceptable range.

While vulnerability scans are essential tools for identifying potential weaknesses in the network, they are not the most critical component for monitoring and securing network endpoints, services, and other vulnerable locations. They provide a snapshot of the vulnerabilities present at a specific point in time, but do not provide continuous monitoring.

SIEM dashboards provide a visual representation of real-time security and network information. However, the dashboard itself does not monitor or secure network endpoints. It merely displays the information gathered by other components of the SIEM system.

Trends are patterns of activity discovered and reported to the SIEM. They help security analysts decide whether reported activity is normal or outside the baseline. However, trends themselves do not monitor or secure network endpoints. They are used to analyze the data collected by other components of the SIEM system.

q_siem_logmgmt_siem_secp8

Which tool or concept used in cybersecurity monitoring gives a condensed overview of information from various data sources for daily incident response tasks?

Answers:

***Security Information and Event Management (SIEM) tools**

Network-based vulnerability scanners

Log files generated by applications on hosts

Host-based intrusion detection systems

Explanation:

SIEM tools aggregate and correlate data from various sources, creating a summarized view of information through event dashboards. These dashboards support daily incident response tasks by providing a comprehensive snapshot of security events and indicators.

Network-based vulnerability scanners focus on identifying vulnerabilities in network devices and systems, but they do not provide summarized views of diverse data sources for incident response tasks.

Although logs created by applications on hosts are sources of data, they are not typically for providing condensed overviews for daily incident response duties.

Host-based intrusion detection systems monitor activities on a single host. They can generate alerts, but are not for aggregating and correlating data from diverse sources for summarized views.

q_mon_metadata_event_metadata_secp8

Which of the following BEST describes the role of event metadata in network security?

Answers:

It provides the specific notification or alert the process raises.

***It is the source and time of the event, which can include a host or network address, a process name, and categorization/priority fields.**

It is used to synchronize each host to the same date, time value, and format.

It is the data that is generated by processes running on network appliances and general computing hosts.

Explanation:

The source and time of the event is the correct answer. Event metadata provides context about the event, including its source and time. This can include a host or network address, a process name, and categorization/priority fields. This information is crucial for understanding and investigating security incidents.

The specific notification or alert the process raises is called event message data, not event metadata.

While synchronization of each host to the same date, time value, and format is important for accurate logging, this is not the role of event metadata.

The data generated by processes running on network appliances and general computing hosts is referred to as event data, which includes both event message data and event metadata.

q_mon_metadata_file_metadata_secp8

What kind of metadata is usually linked with files and includes information like creation date, access history, and security permissions?

Answers:

Web metadata

Email header metadata

Social media metadata

***File metadata**

Explanation:

File metadata encompasses attributes such as creation date, access history, modification details, security permissions, and extended attributes, offering insights into files stored on a system.

Web metadata involves headers exchanged between web clients and servers, encompassing details about the requested resource, data type, and authorization. It does not directly correlate with file attributes like creation date and security permissions.

Email header metadata comprises sender and recipient address details and information about servers engaged in email transmission. It pertains to email communication and does not relate to file attributes.

Social media metadata can encompass details such as the location and time when a user has uploaded media, like photos and videos. This metadata correlates with content uploaded to social media platforms.

q_mon_metadata_ipfix_secp8

You are a cybersecurity analyst tasked with investigating a security incident.

You need to analyze network traffic data that normally goes to Syslog or SNMP, and you also need a tool that can address the need for a standardized protocol for internal protocol flows. Additionally, you need to collect data from routers, servers, and other network appliances.

Which tool would be BEST for you to use?

Answers:

NetFlow

sFlow

***IPfix**

Web metadata

Explanation:

IPfix is the correct answer. IPfix directly integrates data that normally goes to Syslog or SNMP, eliminating additional services collecting data from each network device. IPfix addresses the need for a standardized protocol for internal protocol flows. This data comes from routers, servers, and other network appliances that are mediation systems. The data is formatted, sent to an exporter, and then sent to a collector. IPfix, like NetFlow, looks at flow and the number of packets being sent and received during a given session.

NetFlow is incorrect because while NetFlow is a feature on Cisco routers that can examine each data flow that comes through the network, it does not directly integrate data that normally goes to Syslog or SNMP, nor does it address the need for a standardized protocol for internal protocol flows.

sFlow is incorrect because while sFlow is a packet sampling technology that works on layers 2-7 of the stack, it does not directly integrate data that normally goes to Syslog or SNMP, nor does it address the need for a standardized protocol for internal protocol flows.

Web metadata is incorrect because while web metadata can provide valuable information about user behavior on websites, it is not a tool for analyzing network traffic data or for integrating data that normally goes to Syslog or SNMP. Web metadata includes information like IP addresses, user requests, user downloads, time spent on the site, and attempts to gain unauthorized access, but it does not address the needs specified in the question.

q_mon_metadata_syslog_message_secp8

In the context of a syslog message, which of the following components is calculated from the facility and severity level?

Answers:

Header

Message

***PRI code**

Timestamp

Explanation:

PRI code is the correct answer. The PRI code in a syslog message is calculated from the facility and severity level. The facility refers to the source of the message (like a hardware device, a protocol, or a module of the system software), and the severity level indicates how urgent or critical the message is.

The header of a syslog message contains timestamp, hostname, app name, process ID, and message ID fields. It does not include the PRI code.

The message part of a syslog message contains a tag showing the source process plus content. The format of the content is application-dependent. It does not include the PRI code.

The timestamp in a syslog message indicates the time when the event occurred. It is not calculated from the facility and severity level.

q_mon_metadata_web_metadata_secp8

You are a cybersecurity analyst investigating a potential data breach in your organization. You have identified a suspicious user who appears to have accessed sensitive information.

Which type of metadata would be MOST useful to determining the user's activities on your organization's internal web applications?

Answers:

Email metadata

Mobile metadata

File metadata

***Web metadata**

Explanation:

Web metadata is the correct answer. Web metadata can provide detailed information about a user's activities on web applications, including IP addresses, user requests, downloads, time spent on the site, and attempts to gain unauthorized access. This would be the most useful type of metadata for this investigation.

While email metadata can provide useful information about the sender and recipient of an email, it would not provide detailed information about a user's activities on web applications.

Mobile metadata can provide information about a user's activities on mobile devices, including emails, text messages, and app usage. However, it may not provide detailed information about a user's activities on specific web applications.

File metadata provides information about a file, such as when it was created, accessed, and modified. While this could potentially provide some insight into a user's activities, it would not provide detailed information about a user's activities on web applications.

q_mon_metadata_xheader_secp8

You are worried about email spoofing. What can be put throughout an email's header that provides the originating email account or IP address, and not a spoofed one?

Answers:

***X-headers**

Timestamp

Metadata

Data points

Explanation:

X-headers is the correct choice. You do this with security devices that are designed for this purpose. These devices put X-headers throughout an email's header and provide the originating email account and IP address, not the spoofed one.

Timestamps and metadata (of which timestamps are a part) do not provide an originating email account or IP address (spoofed or otherwise).

A data point is a unique fact such as income, wealth, age of an individual, and the number of dependents. Data points are not designed to provide an originating email account or IP address (spoofed or otherwise).

9.3 Digital Forensics

As you study this section, answer the following questions:

Why is a chain of custody important in an investigation?

What importance does a provable timeline of events play in admissibility of digital forensic evidence?

Why is it important to take a bit-by-bit copy of the logs?

How does the order of volatility help you decide what to secure and preserve first?

What is a digital forensic artifact?

How does provenance play a vital role in digital forensics?

In this section, you will learn to:

Create a forensic drive image with FTK, Guymager, and DC3DD.

Examine a forensic drive image with Autopsy.

The key terms for this section include:

Term	Definition
Legal hold	A process designed to preserve all relevant information when litigation is reasonably expected to occur. A formal notice sent out to all employees of a company when litigation is eminent. The notice instructs all employees to retain electronically stored information (ESI).
Chain of custody	A record of the handling of gathered evidence. This gives all parties involved confidence that no evidence tampering has occurred.
Hashing	A function that converts an arbitrary-length string input to a fixed-length string output.
Provenance	Provenance demonstrates that the digital evidence gathered came from the documented source of evidence and that it has not been tampered with.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.8 Explain appropriate incident response activities. Digital forensics Legal hold Chain of custody Acquisition Reporting Preservation E-discovery 4.9 Given a scenario, use data sources to support an investigation. Data sources Vulnerability scans Dashboards Packet captures

9.3.1 Forensic Documentation and Evidence (Lesson Video)

Transcript:

Digital forensic analysis is the art of revealing important information from computer systems and networks. Every piece of evidence can hold key information. From deleted files to timestamps, user activity, and unauthorized traffic, gathering and examining evidence from computer systems and networks can reveal critical insights. But how do we ensure the validity and integrity of this evidence?

Documentation plays a crucial role in collecting, preserving, and presenting valid digital proofs. Failure to maintain a thorough record can jeopardize the integrity of the evidence. It's important to understand that there are various processes and tools used to acquire different types of digital evidence from computer hosts and networks. By following strict protocols, these processes not only demonstrate how evidence is acquired but also establish that it's an accurate representation of the system at the time of the event.
Just like DNA or fingerprints, digital evidence is latent, meaning it can't be seen with the naked eye. It requires specialized machines and processes to interpret. Therefore, formal steps must be taken to ensure the admissibility of this evidence in court. Proper documentation is crucial in showing how the evidence was collected and analyzed without any tampering or bias.
This is where due process comes into play. Just as the law requires fairness in criminal convictions, forensic investigations must adhere to procedural safeguards to ensure fairness. Anyone involved in the investigation, from technicians to managers, must be aware of these processes to avoid compromising the integrity of the investigation. Defense lawyers will always look for any doubts or mistakes in the evidence or collection process.

When information becomes central to a court case, it's essential to preserve it. This concept, known as legal hold, entails suspending routine deletion of electronic or paper records and logs. Regulatory requirements, industry standards, or litigation notices from law enforcement or lawyers can trigger legal holds, leading to the seizure of computer systems as evidence. When information is subject to legal hold, computer systems may be taken as evidence, disrupting the network.

To maintain the integrity of the evidence, a chain of custody is established. Host devices and media from the crime scene are carefully labeled, bagged, and sealed using tamper-evident bags. Anti-static shielding is used to protect electronic media from damage. Each piece of evidence is documented on a chain of custody form, recording who collected it, who handled it afterward, what measures were taken for backup and analysis, and where it was stored. This ensures the evidence remains intact and untampered with. Finally, the collected evidence is stored in a secure facility. This not only includes access control but also environmental control to protect the electronic systems from damage.

By following these key elements—proper documentation, due process, legal hold, and chain of custody—digital forensic analysts can ensure the validity and credibility of the evidence they gather.

That's it for this lesson. In this lesson, we talked about forensic documentation and evidence. We reviewed the importance of documentation and due process. We discussed legal holds and the importance of the chain of custody for ensuring the integrity of evidence.

9.3.2 Forensic Acquisition of Data (Lesson Video)

Transcript:

A forensic investigator gathers evidence from many sources. In this lesson, we will discuss software and hardware sources of potential evidence. We will also discuss the order in which the evidence needs to be gathered.

The process of capturing volatile data follows the order of volatility. We need to preserve the most volatile data first and work our way down to data that is more persistent. If an attack is underway, your computer forensics response team is probably going to capture volatile data before trying to gather data from the hard disk or turning the system off.

The order of volatility is: RAM, Swap Files/PageFiles, Hard Drive, Remote Logs, Archived Data. These items must be investigated in this order so that potential evidence is not lost.

Each one of these items has its own challenges and potential dangers when acquiring forensic data. Let's look at tools and approaches to acquire and retain data from these items.

In a computer forensics investigation, the data on the hard disk drive is a key piece of evidence. A lot of the things that we do on a computer system are saved in some way on the hard disk drive, including virtual memory. A wealth of data is there.

In addition, information attached to deleted files may still be on the disk. Therefore, the hard drive itself is a goldmine of evidence for a prosecutor. Caution must be exercised when making a copy of the hard drive. A regular file-for-file copy is not good enough. It needs to be a sector-by-sector copy that includes data formerly deleted but still accessible.

RAM is the most volatile of all computer data storage. RAM is cleared when a computer is shutdown. Once gone, it cannot be recovered. Data in RAM can be copied as long as the system is running, but should be done only by someone with proper training. The data stored in RAM can have valuable information. Many times malware like worms, viruses and trojan horses are created as memory-resident only. This makes catching them difficult.

Swap files, or page files, are a virtual extension of RAM. An OS is designed so that if you are running low on RAM it can place not-in-use information into the swap or page file on the hard drive to be used later. An admin can determine how much space to allocate to page files. For the forensic investigator, this is another potential source of evidence. The page file data does not automatically delete at shutdown unless you change the default settings.

The OS is a virtual treasure trove of evidence. File systems like NTFS and Ext4 are the roadmap to where data is stored. There is also a registry which is a database that stores information about all the applications installed on a computer. Registry keys are folder-like objects that store application-specific
information for each application stored in the registry. The recycle bin is where deleted files are placed. Also the print spooler can be examined to see the print history from a computer.

In the last 30 years technology devices have changed dramatically. With the increase of digital devices there is an equal increase in devices that may require forensic examination. These devices contain all the evidence a forensic investigator needs. Each device, such as smart phone, tablet, laptop and smart watch all share common elements, namely RAM, CPU, logs and storage space. A trained investigator must be familiar with all platforms.

Firmware is the basic level of software that controls some hardware. Firmware is stored in read-only, persistent memory. This non-volatile memory does not self-delete and in many cases is not updated. Firmware has become vulnerable to different attacks like rootkits and even steganography. The hard drive has firmware that can be manipulated by malicious actors. Extracting this data is difficult and requires specialized training and expensive tools.
A snapshot can be a useful tool in capturing the exact state of running systems. Snapshots capture volatile data in an ever-changing environment. This tool is not a substitute for the sector-by-sector copy of a disk. The snapshots can be saved for future examination. Snapshots are valuable only if taken of a running system that is currently experiencing an attack.
A cache is stored data that is used to improve the speed of the computing process. There are different caches in a computer system. Address resolution protocol, or ARP, matches IP with MAC addresses. This protocol has a built-in caching component that saves CPU time and network bandwidth by referring to the cache. Internet browser history is also cached and can provide valuable evidence regarding search history.
Network forensic data plays an important role in an investigation. Network appliances that are important in an investigation are firewall, router, switches, domain controllers, DHCP servers, application servers and web proxy servers. There are also applications on the network including intrusion detection and intrusion protection systems.
In computer forensics an artifact is an object that contains forensic value. Forensic value is evidence found on hardware or in areas like registry hives that show indicators of compromise. These indicators are a digital sign of a breach. It is important to preserve these artifacts and understand their value to a forensic investigation.
That's it for this lesson. In this lesson, we learned the procedure used to prioritize evidence gathering called the order of volatility. We discussed different areas within a network that contain potential evidence such as cache, network appliance, OS, disks and RAM. We also covered artifacts and how they are defined and identified.

9.3.3 Forensic Tools (Lesson Video)

Transcript:

In this lesson, we're going to discuss some of the forensic tools on the market. Some of the tools are Linux-specific, and other others are Windows-centric.

Investigating and gathering evidence for court action is part of handling computer security. These computer forensic investigations require specialized tools to gather evidence without making changes to the devices or their stored data. These tools include imaging software, data extraction tools, and advanced search capabilities.
Let's take a look at three of these tools and how they work. We're going to talk about dd-CLI, Memdump, and WinHex.
The first tool, dd, is an extraction tool. It's one of the oldest forensic tools still in use. It's used in Linux and Unix systems to create bit-by-bit copies of a physical hard drive without mounting the drive beforehand. It uses the dd-CLI command, and it's easy to employ. A raw image is created for forensic analysis, and the user has multiple file extension choices.
Memdump is another data gathering tool. It aggregates the data in the volatile RAM memory from devices. Normally, a second tool is needed to search through all the data in a memdump file.
WinHex is a powerful disk and universal hexadecimal editor. It's used in digital forensics for data recovery. WinHex is Windows-compatible and can be used with a whole bunch of options, making it a very versatile forensic tool.
Now let's move on to imaging tools. Imaging tools must create exact copies of data without making a single changeâ€”not even one bit. FTK Imager is a powerful tool that allows the investigator to acquire, preview, and copy data thoroughly and forensically. Autopsy is exactly what it sounds likeâ€”a forensic tool that investigators use to thoroughly examine extracted data, pictures, and empty disk space in order to determine exactly what happened.
And that's it for this lesson. We talked about Forensics tools and how they fall within a few categories; data gathering, imaging, and extraction. We discussed what forensic tools must be able to do--create an exact image of a disk with perfect integrity and then search large amounts of data for specific file extensions, hidden files, deleted data, and root cause analysis.

9.3.4 Create a Forensic Drive Image with FTK (Demo Video)

Transcript:

In this demonstration, we're going to discuss creating a forensic disk image with the Forensic Toolkit. It's better known as the FTK Imager. In our scenario, we need to examine a hard drive from a malicious employee's workstation. Before you try this for real, make sure that your organization considers you qualified to do this and that you have the backing of your legal department.

The first thing we need to do is create an image of the employee's hard drive. Like I said, we're going to use the FTK Imager to do this. Remember, when we're conducting a forensic investigation, we don't want to modify the evidence in any way, shape, or form. If this situation were to be litigated for some reasonâ€""say, the employee gets fired and then sues the company--then the prosecutor could say that when we conducted our forensic investigation, we planted the evidence on the hard drive. That's why it's crucial to make sure that that hard drive isn't changed in any way during the forensics investigation.
We're going to create an image of that hard drive, and we're going to do all of our testing, examining, and investigative work on the imaged copy of the hard drive, not on the actual hard drive. Once we've created that image, we'll use a second tool to analyze the contents of the hard drive to see if we can find anything that's questionable.
The first thing we need to do is get an image of the hard drive. To do this, we have to connect the hard drive to this forensic workstation. That's problematic because as soon as you connect a hard drive to a Windows workstation, Windows immediately starts writing little bits of data to the driveâ€""but we don't want to modify the drive in any way. So we can't just directly connect the drive to a SATA connector on this workstation. Instead, we need to implement a write blocker, also known as a forensic disk controller. Its job is to block writing to the hard drive. It allows us to connect the employee's hard drive to the write blocker, and then we connect the write blocker to our machine, usually with a USB cable. This prevents any write operations coming from the operating system on the forensic computer from going through to the device that we're analyzing. So remember, when you're conducting a forensic investigation on a hard drive and you're going create an image, always use a write blocker to prevent any type of write operation from occurring on that drive. I've already set this system up with a write blocker, and it's ready to go.
Let's go ahead and use the FTK Imager to create an image that we can analyze. By the way, FTK Imager is a free tool. A lot of the forensic software is very, very expensive, but you don't have to spend a lot of money to be able to conduct a good forensic investigation. There are a lot of free, open source, and legitimate tools for forensic investigations.
Within Access Manager, let's click on Add Evidence Item. Then we have to specify what we're going to add. We're adding a physical drive because I've connected it to a write blocker on this particular system. Click Next.
We have to specify which drive we want to create the image from. You'll notice, when I click the dropdown list, it picked all of the hard drives on this system. You need to be very careful that you don't choose the wrong one. For example, if I were to choose Physical Drive 0, that's my local workstation hard disk drive. We'd be creating an image of my local system, not the actual drive that's being used for evidence. Physical Drive 1 is the drive that I've connected to the write blocker. That's what we want to choose. I'll click Finish. We'll come over here, to the Physical Drive 1, and we'll right-click on it. We want to Export Disk Image.
Under Image Destination, we'll click Add. Then we can specify which type of destination image we want to create. Raw (dd) is the default selection because it's probably one of the most widely used imaging format for forensics, so we'll just leave it set the way it is. Click Next.
We need to document the evidence. Remember, whenever you're conducting a forensic investigation, you need to document everything. In fact, you should take a picture of the entire setup that you're using to analyze this hard disk drive. You should take a picture of the drive, take a picture of how it's connected to the write blocker, and take another picture of how the write blocker is connected to the computer. You might even want to take a video of the entire process.
Under Evidence Item Information, you'll want to assign a case number. Let's do â€"Ëœ1 2 3 4' and evidence number â€"Ëœ5 6 7 8'. Then we'll give it a unique description. We'll enter â€"ËœHD from Mary Worley.' I'll put â€"ËœDana Fellows' down as the examiner. Click Next.
We have to specify where we want to store the image file on my computer. I have a folder for my forensic images. Let's create a new folder, â€"Ëœ1 2 3 4 Mary Worley'. I'll put in the same thing here for Image Filename, â€"Ëœ1 2 3 4 Mary Worley'. If we want to, we can fragment the image--that is, break the image file into multiple pieces. This is a small hard drive. I'm going to set that to zero, which basically means we'll have one image file for the entire hard disk. Click Finish.
We want to make sure that Verify the images after they are created is checked. It's also a good idea to create a directory listing of all the files and the image after they're created, so check that box too. It can be useful for when you're searching for information. Click Start. At this point, the imaging process has started. It'll take a little bit of time to complete, especially if you're going to be working with a big hard disk drive. I'm going to pause the recording now and come back when it's done.
Okay. As you can see, it's almost done verifying the image. In just a few seconds, the process will be complete.

Notice that it's created two different hashes, an MD5 Hash and a SHA1 Hash. For both of these, the hashes match. The MD5 Hash matched, and the SHA1 Hash matched as well. That's good. That's exactly what we want to see. I'll go ahead and hit Close.
Before we end this demo, let's verify that the image file has been created. I'll open File Explorer and navigate to the folder where we saved the image. Here are the various pieces of information that were pulled from the hard drive to create it. We have the image file itself. There's a CSV document that contains a listing of all the filenames and directory names. Last on the list, we have a document that provides a nice summary of the image creation process. It gives us the case number, our evidence number, the description, the examiner, all the information we filled out earlier, and the information about the hard drive itself from Mary Worley's computer. At this point, our image file is created. The next step of the process is to use another tool, such as Autopsy, to examine the image.
That's all for this demo. We used FTK imager to capture a forensic image of a hard drive. We discussed the importance of using a write blocker to keep data from getting tampered with and the importance of examining a copy of the disk, not the original disk itself. Then we made a copy the image and verified that it was saved to the folder we created.

9.3.5 Create a Forensic Drive Image with Guymager (Demo Video)

Transcript:

There are several ways to capture a disk image as part of a forensic investigation. In this demo, we're going to do this with a program called Guymager. Guymager has a graphical user interface, making it a bit easier to use than a command line tool. We're on the Guymager home page, and you can read more about it there. One thing I want to point out on the website is that Guymager does come on several live CDs and security operating systems. We're going to use Kali Linux for our demo, so let's close the browser and get started.

I want to mention a couple of things before we get started. Normally, you'll want to have a write blocker between the disk you're imaging and the forensic workstation that you're working from. In a virtual environment, such as my test system here, I can't really do that. The write blocker would keep data from being written to the disk we're wanting to image, which is very important.
The other thing I want to point out is that Kali has a live CD version that has a lot of forensic tools that my copy doesn't have. Keep that in mind if you're setting up your own lab and be sure to investigate the right copy of Linux that will work best for you.
I said that Guymager is a GUI tool, but it does need to be run as a Sudo user. There's a shortcut for it under the Application Launcher, so if I go up here, down to Forensics, and then over to Forensic Imaging Tools, I see Guymager. When it launches, it warns me that it needs to be started with root rights in order to perform acquisitions. That's no fun, and but we can get around it. I'll launch it from the terminal. So let's click No, I don't want to continue here.
I'll go up here and launch a terminal window. After it loads, I'll type â€”Ëœsudo guymager' and press Enter. It prompts me for a password to continue as an admin, so I'll type that in here. Just a reminder, when you type passwords in Linux, the cursor doesn't move, and you think that you're not really typing. This is normal; it's a security feature. Your keyboard isn't broken. I press Enter, Guymager is launched, and now I can acquire images.
I'll make this full screen to take advantage of all the space. Right away, I notice I have three disks, or partitions. This first one is my hard drive for my Kali machine. I know that because it's a 20-gig disk, and that's what I used. This second one is what I'm after. I know this is it because it's a 2-gig disk, which I plugged in. To acquire an image, I simply right-click and choose Acquire Image from the menu.
Now we have a few choices for the file format. This one is called Expert Witness Format. When we select it, we have the option to fill out all this additional data that would be needed if this was part of a legal investigation. We're not going to pick that one. We're going to pick the Linux dd raw image. This is the format I'm going to use when it's time to examine the image content. Over here, we can split the file into smaller pieces. I have a smaller drive, so I could uncheck that box, but I'll just leave it as-is.
Now I need to put the images somewhere. I have a temp folder for these images, and I'll navigate to /home/dana/temp and click Choose to select that directory. I need to supply a filename, so I'll just enter â€”ËœImage1'. The program puts in the extension, so we don't have to do that. Below here, I'm going to make sure this box is checked so that we have the hash value of the image Calculate MD5. We're also going to verify the image after acquisition. Click Start to get the process going.

We have a very small drive, so this shouldn't take long. In fact, if you look at the progress, you can see we're moving along very nicely. It looks like it's finished, and the indicator light is green. We're done with this part, so let's verify that it acquired the images and saved. I'll minimize Guymager for now.
Now let's go up to our Application Launcher and open up File Manager. Remember, we saved the image in a temp directory under /home/dana, and here it is. I'll slide the mouse over and double-click on temp to open it and view the contents. I can see three files. It does look like Guymager broke my file into two, since I checked the box to split files over 2 gigs. The third file is a log file. Let's open up that log and take a look.
My log file has some info about Guymager itself--the version, timestamp, and so on. I'll scroll down here a little ways. Now I see some information about disk size and other data. I want to go down a little farther, to here. This is more relevant information about the acquisition of the disk. I have the device name, /dev/sda, and the size, format, and so on. Down here, I can see the MD5 hash calculation.
Finally, here, we can see when the image was captured, how long it took, and the speed. At the very bottom, we can see the three files that were created when we ran the program. At this point, our disk is captured. The next step is to examine the contents with another forensic software tool designed to do so.
That's it for this demo. In this demo, we used Guymager to capture a drive.

9.3.6 Create a Forensic Drive Image with DC3DD (Demo Video)

Transcript:

In this demonstration, we're going to create a forensic drive image. It's very important that you understand you can't use standard file copying utilities to create a forensic drive image. You can't use, say, Windows Explorer or File Explorer. Those utilities copy files that have an entry, a record, in the allocation table of the partition where it resides, so it only copies data that's associated with a file or folder in the file system. When you're conducting a forensic investigation, you need all the data on that hard drive, especially the data that's not associated with a particular file or file in the file system, but is still on the hard disk drive. Basically, we're looking for stuff that's been deletedâ€""stuff that someone may be trying to hide. We need to use a drive imaging utility to do this. There are a variety of utilities that you can use. Some cost a lot of money; some cost practically nothing. We're going to use the latter option in this demonstration today.

In this demo, we'll use dc3dd to obtain a raw image of a hard drive. dc3dd was developed at the Department of Defense's Cyber Crime Center, and it's basically an enhanced version of the open source dd command with added features for computer forensics. One of the main characteristics of dc3dd is that it offers the possibility of hashing on the fly with multiple algorithms (MD5, SHA-1, SHA-256, or SHA-512). You'll want to use a write blocker between your machine and the disk you're obtaining an image from.
Before we start looking at how to create the drive image, you need to understand how Linux storage devices are addressed by the Linux system. It's kind of difficult to understand when you're new to Linux, but all storage devices on Linux systems are addressed using a device file located in the /dev directory. If a process needs to write information to a hard disk drive, it writes it to a specific file in the /dev directory, which then redirects the IO data to the appropriate hardware device, such as a hard disk drive.
What we need to do is figure out what the device name is for the drive that we want to image. There are a variety of command line utilities you can use to do this. One way to find the forensic image is to use fdisk and the sudo fdisk 1 parameter.
I think we have enough information to get started. We could go to a terminal and start dc3dd, but I'm going to go up to the Application Launcher and start it from there. I'll come down to Forensics and then select Forensic Imaging Tools. I see dc3dd, so let's click on it and launch it.
The only reason I wanted to start it this way is so that you can see we have a nice manual here to learn more about how to use it. I'm not actually going to go over any of this in this information right here, but I will when we type some of the commands in a minute. I'm going to type â€"Ëœclear' and get a clean screen to work with.
I like to see what disks I have to work with by typing â€"Ëœsudo fdisk' with the â€"Ëœ-l' parameter. Press Enter. Now put in the sudo password. The first disk we'll look at is this one, /dev/sdb. It's a 20-gig disk, and I know that this is the one with my Kali Linux installed on. But this isn't the disk I want to image.
If I go up a little, I see a second disk, /dev/sda. This one is 2 gigs. This is the one I want to image. Make sure you know which disk you're working with when you do this. I'll come down here and type â€"Ëœclear' to clear the screen.

Now we need to type in the command to create the image, tell it where to find the disk, where to store the copy, which hash to use, and name the log file. I'll go ahead and type that in and then come back in a second and explain the command.
Okay, I have the command typed in. The first part is â€”Ëœsudo dc3dd'. This just tells Linux to run our command as root, or basically like an admin in Windows. The next part, â€”Ëœif=/dev/sda', is my input file. That's what the letters I-F stand for, input file. After that, we need to specify where our image will go. O-F stands for output file, so we've typed in â€”Ëœof=/home/dana/temp/imaged2.img'. That's the path where I'll save my image file, and I named it â€”Ëœimage2.img'. Next, we'll use â€”Ëœhash=sha256' for our hash type. We could use MD5 or one of the others types we listed in the intro of this demo, but this will work fine. The last thing we have is â€”Ëœlog=/home/dana/temp/image2.log'. As you can guess, this is just a log file, which is very important to have along with the disk image. It's going to be located in the same directory as our disk image. We'll look at after we create it.
This all looks really good. I'll press Enter to start the disk imaging. Down here, we can see the copy progress. Keep in mind that this is a very small disk without a lot of data, so it's going quickly. It looks like it's done. Down here, I can see when it completed. Up here is the location of the image file. I'll go up and close out of our terminal. Now let's go look in our temp file, confirm that it copied over, and look at our log file.
Next, we'll go up to Application Launcher and over to File Manager. I'll go over to my temp folder and double-click on it to open it. Here are my two files, my image file and my log file. Let's open that log now. It shows us things like the time the image started, the size, where the file was located, our hash value, the output filename, and the time it completed. Our next step would be to use another forensic tool to examine the image itself.
And that's it for this demo. We used the command line utility dc3dd to capture a forensic image from a disk.

9.3.7 Examine a Forensic Drive Image with Autopsy (Demo Video)

Transcript:

After a computer forensic investigator captures an image of a drive, they need to examine it. Autospy is a popular examination tool. It's digital forensic software that's free, open source, and said to have most of the features that you'd find in commercial digital forensics tools. Some of the features include, hash matching, registry analysis, and web analytics, and the ability to do a keyword search. For our demo, we're going to use Autopsy to analyze a disk that has been previously captured and saved to my hard drive.

Double-click on it to start it. Sometimes it takes a minute or so for it to fully launch. We want to create a new case. Click on that option. We'll give it a case name, â€”Ëœ1 2 3 4 Mary Worley'. Let's specify the directory where we're going to store the information we're about to create. Let's put it in the same directory as our image file. We'll keep all the data together. I'll actually create a folder for it and name it â€”ËœAutopsy'. I'll select the folder we just created and then click Next. Enter in a case number, â€”Ëœ1 2 3 4 Mary Worley'. I'll put the examiner's name in here. I'll hit Finish.
This next part is going to take a minute or two while it creates the case and gathers all the files. I'll pause the demo while this is running.
Okay, that took several minutes, and now I'm presented with the Add Data Source page. We have to specify what we want to analyze. We're going to analyze a disk image or VM file. Click Next. We need to specify what image file we want to look at. It's in a folder called Forensic Images under 1 2 3 4 Mary Worley. There's the image file that we created a minute ago. Hit Open. We need to set our time zone. I'm in Mountain Time, so I'll scroll down until I find that and select it. Click Next.
We have to specify exactly what we want to look for. This tool uses what's called Ingest Modules. An ingest module is basically just a piece of software that looks for a particular type of information in the image file. You can come over here and specify what you want to look for. For example, you could go under Keyword Search and specify what type of information you want that particular module to look for. We want to look for all of this information here, so I'll check all the boxes and click Next. Now let's go ahead and click Finish.
The process of analyzing the image file has started. It can take quite some time to complete, depending on the size of disk and amount of data. In my test environment, I don't have a very large disk or very much data, so it should happen pretty quickly. But I'll go ahead and pause the demo while it runs.
At this point, the image file has been thoroughly examined by the Autopsy tool, and the results are displayed here on the left, organized by the type of data or the view of the data that you want to use. It's very useful. For example, it sorts the data by the file type extension.

Let's click on Images. There are 43 of them. Let's see what we have. If I click this one, WACC classroom.jpg, I can see the image down here, in a preview pane. It looks like a picture of a computer lab or classroom.
See this one up here, with the red X? This is a file that was deleted on the hard drive. Of course, as you probably know, when you delete something, you don't actually delete it at all--you just delete the pointer, or reference to it. The file won't go away until the drive is formatted or overwritten. Even then, files are sometimes recoverable. Down in the preview, this looks like a screen shot of something.
I'll click on Videos. Over in the listing, I have four videos. One's been deleted. When I click on Audio, my list is the same as it was for videos, so Autopsy must recognize that there are audio files as part of the videos. Under Documents, I have some Microsoft Office documents. I could export those out and look at the contents in the supported programs if I wanted to. I have a few PDF documents as well, along with some plain text files. Once again, a few of these have been deleted, as you can see by the red X. I have no executable files to look at, so we'll skip those. Autopsy does give us a list of just the deleted files, all grouped together. That is handy if I was focused just on what Mary may have deleted and might be trying to hide.
I have a few more categories. We have Recycle Bin and something called User Content Suspected. I'll make my way down to the email addresses. Out of concern for confidentiality, I'm not going to open this up because I'm not sure exactly what's on this disk. Just to clarify, when we see things like emails addresses, IP addresses, URLs, and so on, it does not mean email accounts configured on the system, but any email that shows up on the disk in a document, spreadsheet, etc.
Under IP addresses, I see 0.1.2.3. This might be a simple false positive because the format is somewhat like an IP address. Under phone numbers, I have some phony-looking phone numbers. This looks like something from TV, since they have a 555 prefix. I have a bunch of URLs listed. This is handy for seeing where the person has been spending their time on the web.
I have Hashset Hits, but there's nothing there to see. Right under this, I actually have Email Messages, but there aren't any to look at on this disk. I have Credit Cards down here. Once again, I'm not going to click on those and open them because I'm not sure what I might find. And now we're at the bottom of the list.
Now, there are entire courses and degree programs on how to use this product and the whole legal process that goes along with it. This demo is a very brief overview of what the software is capable of doing. We didn't even look at things such as geolocation, timeline, report generation, or other featuresâ€”there's a lot more to learn!
That's it for this demonstration. In this demo, we used Autopsy to examine a disk image.

9.3.8 Forensic Data Integrity and Preservation (Lesson Video)

Transcript:

Data integrity and preservation are fundamental aspects of digital forensics, ensuring the reliability and usefulness of the data. They serve to maintain and protect the original state of digital evidence from the moment of collection to its presentation in court. Any intentional or unintentional alteration can compromise the accuracy of the evidence, thereby undermining the entire forensic investigation. This means that evidence acquisition, imaging, and storage all need to be done using well-defined processes.

Evidence should be captured in the order of volatility, from more volatile to less volatile. The ISOC best practice guide to evidence collection and archiving recommends the following order: First, CPU registers and cache memory—including cache on disk controllers, graphics cards, etc. The contents of nonpersistent system memory, or RAM, including routing table, ARP cache, process table, and kernel statistics. Then, data is stored on persistent mass storage devices like HDDs, SSDs, and flash memory devices, followed by partition and file system blocks, slack space, and free space. Next would be system memory caches, such as swap space, virtual memory, and hibernation files. And finally, user, application, and OS files and directories. Other sources could include remote logging and monitoring data, physical configuration and network topology, and archival media or printed documents.
Data acquisition is complicated by the fact that it's more difficult to capture evidence from a digital crime scene than it is from a physical one. For example, some evidence could be lost if a computer system is powered off; on the other hand, other evidence may be unobtainable until after the system has been powered off. Because of this, it's important to know the three states for persistent storage acquisition.
Live acquisition means copying the data while the host is still running. This may capture more evidence or more data for analysis and reduce the impact on overall services. Still, the data on the actual disks will have changed, so this method

may not produce legally acceptable evidence. It may also alert the threat actor and allow time for them to perform anti-forensics.
With static acquisition, the host is shut down normally. If malware is a concern, this method runs the risk that the malware will detect the shutdown process and perform anti-forensics to try to remove traces of itself.
If this is the case, static acquisition by pulling the plug instead of powering down may be a solution. This means disconnecting the power at the wall socket (not the hardware power-off button). This is most likely to preserve the storage devices in a forensically clean state, but there's the risk of corrupting data.
Once the target disk has been safely attached to the forensics workstation, data acquisition proceeds as follows:
A cryptographic hash of the source disk media is made using either the MD5 or SHA hashing function. A bit-by-bit copy of the source media is made using an imaging utility. A second hash is then made of the image, which should match the original hash of the media. A copy is made of the reference image, validated again by the checksum. All analysis is performed on the copy.
This proof of integrity ensures non-repudiation. If the provenance of the evidence is certain, the threat actor identified by analysis of the evidence can't deny their actions. The hashes prove that no modification has been made to the image.
That's it for this lesson. In this lesson, we discussed data integrity and preservation. We reviewed the order of volatility that requires that evidence be captured in order of volatility from more volatile to less volatile. We also talked about ensuring acquisition integrity by leveraging the different states of persistent storage acquisition: live acquisition, static acquisition, powered down, and static acquisition, unplugged. Lastly, we discussed how proof of evidence integrity helps to ensure non-repudiation.

9.3.9 Forensic Investigation Facts

Digital forensic analysis involves examining evidence gathered from computer systems and networks to uncover relevant information, such as deleted files, timestamps, user activity, and unauthorized traffic. There are many processes and tools for acquiring different kinds of digital evidence from computer hosts and networks. These processes must demonstrate exactly how the evidence was acquired and that it is a true copy of the system state at the time of the event.

This lesson covers the following topics:

- Due process and legal hold
- Chain of custody
- Reporting
- Data sources
- System memory acquisition
- Data image acquisition

Due Process and Legal Hold

Digital forensics is the practice of collecting evidence from computer systems to a standard that will be accepted in a court of law. Forensics investigations are most likely to be launched to prosecute crimes arising from insider threats, notably fraud or misuse of equipment. Prosecuting external threat sources is often difficult, as the threat actor may well be in a different country or have taken effective steps to disguise their location and identity. Such prosecutions are normally initiated by law enforcement agencies, where the threat is directed against military or governmental agencies or is linked to organized crime.

Due process is a term used in US and UK common law to require that people only be convicted of crimes following the fair application of the laws of the land. More generally, due process can be understood to mean having a set of procedural safeguards to ensure fairness. This principle is central to forensic investigation. If a forensic investigation is launched (or if one is a possibility), it is important that technicians and managers are aware of the processes that the investigation will use. It

is vital that they can assist the investigator and that they do not do anything to compromise the investigation. In a trial, defense counsel will try to exploit any uncertainty or mistake regarding the integrity of evidence or the process of collecting it.

Legal hold refers to the fact that information that may be relevant to a court case must be preserved. Information subject to legal hold might be defined by regulators or industry best practice, or there may be a litigation notice from law enforcement or lawyers pursuing a civil action. This means that computer systems may be taken as evidence, with all the obvious disruption to a network that entails. A company subject to legal hold will usually have to suspend any routine deletion/destruction of electronic or paper records and logs.

Chain of Custody

The host devices and media taken from the crime scene should be labeled, bagged, and sealed, using tamper-evident bags. It is also appropriate to ensure that the bags have antistatic shielding to reduce the possibility that data will be damaged or corrupted on the electronic media by electrostatic discharge (ESD). Each piece of evidence should be documented by a chain of custody form. The chain of custody documentation records where, when, and who collected the evidence, who subsequently handled it, and where it was stored. This establishes integrity and proper handling of evidence. When security breaches go to trial, the chain of custody protects an organization against accusations that evidence has either been tampered with or is different than it was when it was collected. Every person in the chain who handles evidence must log the methods and tools they used.

The evidence should be stored in a secure facility; this not only means access control, but also environmental control, so that the electronic systems are not damaged by condensation, ESD, fire, and other hazards.

Reporting

Digital forensics reporting summarizes the significant contents of the digital data and the conclusions from the investigator's analysis. It is important to note that strong ethical principles must guide forensics analysis:

- Analysis must be performed without bias. Conclusions and opinions should be formed only from the direct evidence under analysis.
- Analysis methods must be repeatable by third parties with access to the same evidence.
- Ideally, the evidence must not be changed or manipulated. If a device used as evidence must be manipulated to facilitate analysis (disabling the lock feature of a mobile phone or preventing a remote wipe, for example), the reasons for doing so must be sound and the process of doing so must be recorded.

Defense counsel may try to use any deviation of good ethical and professional behavior to have the forensics investigator's findings dismissed.

A forensic examination of a device that contains electronically stored information (ESI) entails a search of the whole drive, including both allocated and unallocated sectors, for instance. E-discovery is a means of filtering the relevant evidence produced from all the data gathered by a forensic examination and storing it in a database in a format such that it can be used as evidence in a trial. E-discovery software tools have been produced to assist this process. Some of the functions of e-discovery suites are as follows:

- Identify and de-duplicate files and metadata —many files on a computer system are "standard" installed files or copies of the same file. E-discovery filters these types of files, reducing the volume of data that must be analyzed.
- Search —allow investigators to locate files of interest to the case. As well as keyword search, software might support semantic search. Semantic search matches keywords if they correspond to a particular context.

Tags —apply standardized keywords or labels to files and metadata to help organize the evidence. Tags might be used to indicate relevancy to the case or part of the case or to show confidentiality, for instance.

Security —at all points, evidence must be shown to have been stored, transmitted, and analyzed without tampering.

Disclosure —an important part of trial procedure is that the same evidence be made available to both plaintiff and defendant. E-discovery can fulfill this requirement. Recent court cases have required parties to a court case to provide searchable ESI rather than paper records.

Data Sources

The following table outlines several data sources for forensic investigations along with their descriptions.

Source	Description
Dashboards	An event dashboard provides a console to work from for day-to-day incident response. It provides a summary of information drawn from the underlying data sources to support some work tasks. Separate dashboards can be created to suit many different purposes. An incident handler's dashboard will contain uncategorized events that have been assigned to their account, plus visualizations (graphs and tables) showing key status metrics. A manager's dashboard would show overall status indicators, such as number of unclassified events for all event handlers.
Log data	Log data is a critical resource for investigating security incidents. As well as the log format, you must also consider the range of sources for log files, and know how to determine what type of log file will best support any given investigation scenario.
Host operating system logs	An operating system (OS) keeps a variety of logs to record events as users and software interact with the system. Different log files represent different aspects of system functionality. These files are intended to hold events of the same general nature. Some files hold events from different process sources; others are utilized by a single source only.
Linux logs	Linux logging can be implemented differently for each distribution. Some distibutions use syslog to direct messages relating to a particular subsystem to a flat text file. Other distributions use Journald as a unified logging system with a binary, rather than plaintext, file format. Journald messages are read using the journalctl command, but it can be configured to export some messages to text files via syslog.
Windows logs	The three main Windows event log files are the following: Application—events generated by application processes, such as when there is a crash, or when an app is installed or removed. Security—audit events, such as a failed login or access to a file being denied. System—events generated by the operating system's kernel processes and services, such as when a service or driver cannot start, when a service's startup type is changed, or when the computer shuts down.
Application logs	An application log file is simply one that is managed by an application rather than the OS. The application may use Event Viewer or syslog to write event data using a standard format, or it might write log files to its own application directories in whatever format the developer has selected.

Source	Description
	In Windows Event Viewer, there is a specific application log, which can be written to by any authenticated account. There are also separate custom application and service logs, which are managed by specific processes. The app developer chooses which log to use, or whether to implement a logging system without using Event Viewer. Check the product documentation to find out where events for a particular software app are logged.
Endpoint logs	An endpoint log is likely to refer to events monitored by security software running on the host, rather than by the OS itself. This can include host-based firewalls and intrusion detection, vulnerability scanners, and antivirus/antimalware protection suites. Suites that integrate these functions into a single product are often referred to as an endpoint protection platform (EPP), enhanced detection and response (EDR), or extended detection and response (XDR). These security tools can be directly integrated with a SIEM using agent-based software.
Network logs	Network logs are generated by appliances such as routers, firewalls, switches, and access points. Log files will record the operation and status of the appliance itself—the system log for the appliance—plus traffic and access logs recording network behavior.
IPS/IDS logs	An IPS/IDS log is an event when a traffic pattern is matched to a rule. As this can generate a very high volume of events, it might be appropriate to only log high sensitivity/impact rules. As with firewall logging, a single packet might trigger multiple rules.
File	File metadata is stored as attributes. The file system tracks when a file was created, accessed, and modified. A file might be assigned a security attribute, such as marking it as read-only or as a hidden or system file. The ACL attached to a file showing its permissions represents another type of attribute. Finally, the file may have extended attributes recording an author, copyright information, or tags for indexing/searching.
Web	When a client requests a resource from a web server, the server returns the resource plus headers setting or describing its properties. Also, the client can include headers in its request. One key use of headers is to transmit authorization information, in the form of cookies. Headers describing the type of data returned (text or binary, for instance) can also be of interest. The contents of headers can be inspected using the standard tools built into web browsers. Header information may also be logged by a web server.
Email	An email's Internet header contains address information for the recipient and sender, plus details of the servers handling transmission of the message between them. When an email is created, the mail user agent (MUA) creates an initial header and forwards the message to a mail delivery agent (MDA). The MDA should perform checks that the sender is authorized to issue messages from the domain. Assuming the email isn't being delivered locally at the same domain, the MDA adds or amends its own header and then transmits the message to a message transfer agent (MTA). The MTA routes the message to the recipient, with the message passing via one or more additional MTAs, such as SMTP servers operated by ISPs or mail security gateways. Each MTA adds information to the header. Headers aren't exposed to the user by most email applications. You can view and copy headers from a mail client via a message properties/options/source command. MTAs can add a lot of information in each received header, such as the results of spam checking. If you use a plaintext editor to view the header, it can be difficult to identify where each part begins and ends. Fortunately, there are plenty of tools available to parse headers and display them in a more structured format. One example is the Message Analyzer tool, available as part of the Microsoft Remote Connectivity Analyzer (testconnectivity.microsoft.com/tests/o365). This will lay out the hops that the message took more clearly and break out the headers added by each MTA.

System Memory Acquisition

System memory is volatile data held in Random Access Memory (RAM) modules. Volatile means that the data is lost when power is removed. A system memory dump creates an image file that can be analyzed to identify the processes that are running, the contents of temporary file systems, registry data, network connections, cryptographic keys, and more. It can also be a means of accessing data that is encrypted when stored on a mass storage device.

A specialist hardware or software tool can capture the contents of memory while the host is running. Unfortunately, this type of tool needs to be preinstalled as it requires a kernel mode driver to dump any data of interest. Various commercial tools are available to perform system memory acquisition on Windows.

Data Image Acquisition

Disk image acquisition refers to acquiring data from nonvolatile storage. Nonvolatile storage includes hard disk drives (HDDs), solid state drives (SSDs), firmware, other types of flash memory (USB thumb drives and memory cards), and optical media (CD, DVD, and Blu-ray). This can also be referred to as device acquisition, meaning the SSD storage in a smartphone or media player. Disk acquisition will also capture the OS installation if the boot volume is included.

Given sufficient time at the scene, an investigator might decide to perform both a live and static acquisition. Whichever method is used, it is imperative to document the steps taken and supply a timeline and video-recorded evidence of actions taken to acquire the evidence.

It is vital that the evidence collected at the crime scene conforms to a valid timeline. Digital information is susceptible to tampering, so access to the evidence must be tightly controlled. Video recording the whole process of evidence acquisition establishes the provenance of the evidence as deriving directly from the crime scene.

To obtain a forensically sound image from nonvolatile storage, the capture tool must not alter data or metadata (properties) on the source disk or file system. Data acquisition would normally proceed by attaching the target device to a forensics workstation or field capture device equipped with a write blocker. A write blocker prevents any data on the disk or volume from being changed by filtering write commands at the driver and OS level.

9.3.10 Practice Questions (Section Quiz)

q_for_invest_chain_01_secp8

What is the MOST important element related to evidence in addition to the evidence itself?

Answers:

Photographs of the crime scene

***Chain of custody document**

Witness testimony

Completeness

Explanation:

The chain of custody document is the most important item related to the evidence.

Nothing is more important than the chain of custody document, including photographs.

Witness testimony can be helpful, but it is not more important than the chain of custody document.

Completeness of the evidence is beneficial, but it is not as beneficial as a reliable chain of custody document.

q_for_invest_chain_02_secp8

What is the purpose of a chain of custody?

Answers:

Retaining evidence integrity.

Detailing the timeline between creation and discovery of evidence.

Identifying the owner of the evidence.

***Arriving at conclusions from the investigator's analysis.**

Explanation:

The chain of custody is used to track the people who came in contact with the evidence. The chain of custody starts at the moment evidence is discovered and lists the identity of the person who discovered, logged, gathered, protected, transported, stored, and presented the evidence. The chain of custody helps to ensure the admissibility of evidence in court.

Legal hold refers to the fact that information that may be relevant to a court case must be preserved. This includes retaining (preserving) the integrity of evidence.

Given sufficient time at the scene, an investigator might decide to perform both a live and static acquisition. As part of this process, the investigator normally details the timeline between the creation and discovery of evidence.

Digital forensics reporting summarizes the significant contents of the digital data and the conclusions from the investigator's analysis.

q_for_invest_chain_03_secp8

You have been asked to draft a document related to evidence-gathering that contains details about personnel in possession and control of evidence from the time of discovery up through the time of presentation in court.

Which type of document is this?

Answers:

CPS (certificate practice statement)

FIPS-140

Rules of evidence

***Chain of custody**

Explanation:

The chain of custody is a document related to evidence-gathering that contains details about personnel in possession and control of evidence from the time of discovery up through the time of presentation in court.

A CPS (certificate practice statement) is a document written by a certificate authority outlining their certificate handling, management, and administration procedures.

FIPS-140 is a government standard that defines procedures, hardware, and software that can be employed when performing forensic investigations of cybercrime.

The rules of evidence are the restrictions that must be adhered to in order to ensure the admissibility of collected evidence.

q_for_invest_dashboards_secp8

A CEO asks the tech department to create a console that shows day-to-day incident response and summaries of information drawn from underlying data sources.

What can the tech department present to the CEO as a viable option?

Answers:

***Dashboards**

Network logs

Metadata

Log data

Explanation:

An event dashboard provides a console to work from for day-to-day incident response and a summary of information drawn from the underlying data sources to support some work tasks.

Generated by appliances such as routers, firewalls, switches, and access points, network logs will record the operation and status of the appliance itself, the system logs for the appliance, plus traffic and access logs recording network behavior.

Metadata is the properties of data created by an application, stored on media, or transmitted over a network.

Log data is a critical resource for investigating security incidents. Although kept and analyzed on each host individually, most organizations require better visibility into data sources and host monitoring.

q_for_invest_data_image_secp8

You are a digital forensic analyst and you need to acquire a data image from a computer system as part of an ongoing investigation.

Which of the following steps should you take to ensure the integrity of the evidence and the investigation?

Answers:

Turn off the computer system immediately to prevent any further changes to the data.

***Use a write-blocking device when creating the data image to prevent any changes to the original data.**

Make changes to the system settings to facilitate the data image acquisition process.

Keep documenting the process of data image acquisition, including the tools and methods used, as a low priority.

Explanation:

Using a write-blocking device when creating the data image ensures that the original data cannot be altered during the acquisition process is the correct answer. This is a crucial step in preserving the integrity of the evidence.

Turning off the computer system immediately could result in loss of volatile data and could potentially alter the state of the system, which could compromise the integrity of the evidence.

Making changes to the system settings could alter the state of the system and the data, which could compromise the integrity of the evidence. Any changes to the system should be avoided during the data image acquisition process.

Documenting the process of data image acquisition is a high priority, including the tools and methods used, is a crucial part of the investigation. This documentation can serve as a record of the steps taken during the investigation and can be used to verify the integrity of the evidence and the investigation process.

q_for_invest_data_sources_secp8

You are a digital forensic analyst working on a high-profile case.

You have been given access to a variety of data sources, including dashboards, log data, and host operating system logs. You need to determine the most effective way to gather evidence for your investigation.

Which of the following approaches would be the MOST effective?

Answers:

Rely solely on the dashboard as it provides a summary of information drawn from the underlying data sources.

Focus only on the log data, as it is a critical resource for investigating security incidents.

Concentrate on the host operating system logs as they record events as users and software interact with the system.

***Utilize all the data sources (dashboards, log data, and host operating system logs) to gather a comprehensive set of evidence.**

Explanation:

The most effective approach would be to utilize all the data sources, such as dashboards, log data, and host operating system logs. Each of these sources provides a different type of information and can contribute to a comprehensive understanding of the events that occurred. By using all available data sources, a digital forensic analyst can gather the most complete set of evidence for their investigation.

Relying solely on the dashboard would not be the most effective approach. While dashboards provide a summary of information, they may not contain all the detailed data needed for a thorough investigation. They are best used for day-to-day incident response and to provide an overview of the situation.

Focusing only on the log data would also not be the most effective approach. While log data is indeed a critical resource for investigating security incidents, it is just one piece of the puzzle. It provides valuable information about events that have occurred, but it may not provide the full context needed for a comprehensive investigation.

Concentrating only on the host operating system logs would not be the most effective approach either. These logs do record events as users and software interact with the system, but they represent just one aspect of system functionality. They may not provide a complete picture of the events that occurred.

q_for_invest_due_process_secp8

In cybersecurity investigations, why is it crucial to ensure the admissibility of digital evidence collected from computer systems?

Answers:

Digital evidence is often visible to the naked eye, ensuring its authenticity.

***Due process and the fair application of laws require proper handling of digital evidence.**

The location and identity of threat actors are easily identifiable through digital evidence.

Threat actors can tamper with digital evidence without affecting its integrity.

Explanation:

Due process demands that incident response teams handle evidence collected in cybersecurity investigations meticulously and fairly. Due process ensures that individuals are only convicted of crimes based on unbiased and admissible evidence, contributing to the integrity of the legal system.

Digital evidence, like DNA or fingerprints, is latent, meaning the naked eye cannot see it, and it requires proper interpretation using specialized processes or machines.

Digital forensics helps reveal the location and identity of threat actors through a systematic investigative approach, even if they have hidden their information.

Preserving digital evidence without changes is crucial for maintaining its authenticity and reliability. Proper documentation and forensic procedures are necessary for admissibility in court and trustworthy investigations.

q_for_invest_e-discovery_secp8

A lawyer is preparing a subpoena for an upcoming cybercrime case and is consulting with a digital forensics specialist.

The lawyer explains the need for the ability to parse through data quickly and provide a copy of everything found to the opposing counsel.

Which utility can accomplish these requests?

Answers:

***E-discovery**

Legal hold

Due process

Live acquisition

Explanation:

E-discovery is the means of gathering digital information for a court case. Several e-discovery utilities provide searches based on syntax, keywords, and file type. The lawyer could easily share this type of electronically stored information with opposing counsel.

Legal hold refers to the maintainance of information potentially relevant to a case. Legal holds include taking papers, hard drives, CDs, workstations, and servers.

Due process is the legal term for a a fair trial that leads to being sentenced for a crime or admonished of it. Procedural safeguards are in effect to ensure fairness throughout the trial.

Live acquisition means copying information while the system is still running. Information gathered via this means is not always admissible.

q_for_invest_evidence_01_secp8

A company's cybersecurity team investigates a potential data breach on its network by examining various data sources to gather evidence.

What is the significance of ensuring the evidence is a true copy of the system state at the time of the breach?

Answers:

It prevents unauthorized user access to the investigation.

It accelerates recovery by restoring the system to the state before the breach.

***It establishes the credibility and integrity of the evidence.**

It eliminates the need for documentation in the investigation.

Explanation:

It is essential to ensure that the evidence accurately represents the system state during the breach to maintain its credibility and integrity. This process helps validate the authenticity of the evidence and aids in a successful investigation.

Although ensuring a true copy of the system state does not primarily prevent unauthorized access to the investigation, it is crucial for maintaining the accuracy and reliability of the evidence.

While restoring the system might be important for recovery, it does not address the importance of accurate evidence for legal proceedings.

Evidence is always accompanied by the proper documentation to help verify the existence and validity of the evidence.

q_for_invest_evidence_02_secp8

While investigating a potential cybercrime, a junior digital forensics specialist leaves an important hard drive in a public area overnight.

A senior digital forensics specialist finds the hard drive in the morning and says that it is no longer evidence in the case.

What made the hard drive unusable in court? (Select two.)

Answers:

***The forensics team did not maintain the chain of custody.**

***The forensics team did not maintain the provenance of the hard drive.**

The forensics team did not maintain the legal hold of the hard drive.

The forensics team did not maintain the order of volatility for the hard drive.

The forensics team did not provide a digital forensics report.

Explanation:

The following aspects of this situation make the hard drive unusable in court:

Maintaining the proper chain of custody on any physical evidence at all times is crucial. The hard drive left overnight in a public location is suspect for compromise, rendering it inadmissible in a court of law.

Provenance shows that evidence moves directly from the crime scene without tampering. Provenance is no longer provable, as the hard drive was left alone overnight.

Legal hold refers to preserving information relevant to a court case. Legal hold includes data or files, both paper and digital. It does not impact the usability of evidence.

The order of volatility is capturing data concerning volatility from most to least volatile. Order of volatility would not make a hard drive inadmissible in court. It does not impact the usability of evidence.

Digital forensics reporting summarizes the significant contents of the digital data and the conclusions from the investigator's analysis. It does not impact the usability of evidence.

q_for_invest_file_metadata_secp8

A forensic analyst at an international law enforcement agency investigates a sophisticated cyber-espionage case.

The analyst must uncover the timeline of document interactions, detect concealed or system-protected files, interpret categories of digital events, and trace digital breadcrumbs left behind during media uploads on social platforms.

What combination of data sources would provide the MOST comprehensive information for this multifaceted investigation?

Answers:

***File metadata with extended attributes and network transaction logs.**

Network transaction logs and gateway security logs.

File metadata and event logs.

Event logs and gateway security logs.

Explanation:

Paired with network transaction logs, file metadata with extended attributes provides a comprehensive understanding of document interactions, including hidden details from online actions and network operations.

Network transaction logs and gateway security logs offer insights into network operations and rule-based activities, but need more depth for understanding file interactions or hidden information in online uploads.

File metadata and event logs provide a detailed view of file creation, access, modifications, and event sources, but have limitations in capturing online media upload traces.

Event logs and gateway security logs mainly focus on event sources and network connections, missing the in-depth view of document-specific details or hidden clues from online interactions.

q_for_invest_legal_hold_01_secp8

Your company is about to begin litigation, and you need to gather information. You need to get emails, memos, invoices, and other electronic documents from employees. You'd also like to get printed, physical copies of documents.

Which tool would you use to gather this information?

Answers:

***Legal hold**

Chain of custody

Timestamps

Timeline of events

Explanation:

You would use a legal hold. The purpose behind a legal hold is to help ease the burden of the IT and legal teams as they gather evidentiary documentation. This notice instructs employees to retain any electronically stored information, or ESI.

The chain of custody proves that no one has tampered with the gathered evidence .

Timestamps provide an exact date and time of an event and must be accurate to be admissible.

A timeline of events is required for digital forensic evidence to be admissible and to prove who is most responsible for what occurred.

q_for_invest_legal_hold_02_secp8

A cybersecurity manager is preparing to begin working when a police officer comes through the door waving a subpoena. The officer states that the company is under investigation for suspicious activities relating to recent overseas sales, and they are taking the servers with them.

What gives police officers the right to take the servers?

Answers:

***Legal hold**

Due process

Digital forensics

Data acquisition

Explanation:

Legal hold refers to the maintenance of information potentially relevant to a case. Legal holds include taking papers, hard drives, CDs, workstations, and servers.

Due process is the legal term for a fair trial before someone is sentenced for a crime or absolved. Procedural safeguards are in effect to ensure fairness throughout the trial.

Digital forensics is the process of obtaining the information required for a court case. It also determines procedures that ensure the information is admissible in a court of law.

Data acquisition is the tools and techniques used to create a forensically sound copy of the data. The data cannot have changes to be admissible.

q_for_invest_memory_dump_secp8

As a digital forensics investigator, you are tasked with investigating a potential data breach in your organization. You suspect that a sophisticated malware has infiltrated the system and is deleting its traces from the hard drive after executing its operations.

Which of the following steps would be the MOST effective in capturing the evidence of this malware's activity?

Answers:

Running a full system antivirus scan.

Checking the system's event logs.

***Performing a system memory dump.**

Conducting a network traffic analysis.

Explanation:

Performing a system memory dump is the correct answer. A system memory dump involves capturing the contents of the system's RAM. Since the suspected malware is running and then deleting its traces, its activities would be present in the system memory. A memory dump would provide a snapshot of the system's state at a particular point in time, including the activities of all running processes, which would provide the necessary evidence.

While running a full system antivirus scan is a standard procedure in investigating potential malware infections, sophisticated malware can often evade antivirus scans. Moreover, if the malware is deleting its traces from the hard drive, the antivirus scan might not find any evidence of its activity.

System logs can provide valuable information about system activities, including potential malware operations. However, advanced malware often has the capability to delete or alter these logs to cover its tracks. Therefore, relying solely on system logs might not provide the necessary evidence.

While network traffic analysis can provide evidence of malware communication with external servers, it might not provide direct evidence of the malware's activity on the system. Furthermore, if the malware is using advanced techniques like encryption or steganography to hide its network communication, the network traffic analysis might not reveal any suspicious activity.

q_for_invest_recording_secp8

When conducting a cybersecurity investigation, how does recording the evidence acquisition process on video help to ensure the collected evidence's integrity?

Answers:

Video recording ensures that no one can tamper with the evidence.

Video recording provides a backup of the collected digital data.

Video recording verifies the authenticity of the forensic workstation.

***Video recording proves the evidence originated directly from the crime scene.**

Explanation:

Video recording of the evidence acquisition process establishes its provenance, demonstrating that law enforcement retrieved it directly from the crime scene. This visual record helps validate the authenticity and integrity of the collected evidence.

While video recording can help establish the integrity of the evidence-collection process, it does not guarantee that someone cannot tamper with the evidence after collection.

Video recording focuses on capturing the evidence acquisition process rather than providing a backup of the collected digital data.

While video recording can provide insights into the activities of the forensic workstation, its primary purpose is to document the evidence acquisition process and its connection to the crime scene.

q_for_invest_reporting_secp8

As a digital forensic analyst, you have completed an investigation and are now tasked with creating a report summarizing your findings.

Which of the following principles should guide your report writing?

Answers:

The report should be biased towards the hypothesis you initially formed about the case.

The analysis methods used should not be repeatable by third parties.

***The evidence must not be changed or manipulated unless necessary. If it is changed or manipulated, the reasons why and process used must be recorded.**

The report should only include conclusions and opinions formed from the direct evidence under analysis.

Explanation:

The evidence in a digital forensics investigation must not be changed or manipulated unless it is necessary for the analysis. If the evidence must be manipulated, the reasons for doing so and the process of doing so must be recorded. Recording this information ensures the integrity of the evidence and the investigation process.

A digital forensics report should not be biased toward any hypothesis. The report should be objective and based on the evidence found during the investigation. Any bias could compromise the integrity of the investigation and the report.

The analysis methods used in a digital forensics investigation should be repeatable by third parties. Repeatable results ensure that the investigation process is transparent and can be verified by others if needed.

The report should only include conclusions and opinions that are formed from the direct evidence under analysis. Following this policy ensures that the report is based on factual information and not on speculation or assumptions.

9.4 Redundancy

As you study this section, answer the following questions:

Why is redundancy important to network security?

Why would an organization use geographic dispersal?

What are the levels of RAID and when would you use each level?

Why would a system administrator want to use load balancers?

What is an uninterrupted power supply used for?

What is the difference between active/active and active/passive?

What is the main advantage of RAID 0? Disadvantage?

What is the difference between RAID 0+1 and RAID 1+0?

Key terms for this section include the following:

Term	Definition
Fault tolerance	The ability to respond to an unexpected hardware or software failure without loss of data or loss of operation.
Redundancy	A method for providing fault tolerance by using duplicate or multiple components that perform the same function.
Geographic dispersion	Using multiple locations to store data to mitigate downtime due to loss of availability at a location.

Multipath	A fault-tolerance technique that gives multiple physical paths between a CPU and a mass-storage appliance.
Load balancers	A process that distributes processing among multiple nodes.
Uninterrupted power supply (UPS)	A stand-alone power supply that allows servers to be gracefully shutdown during a power outage.
Virtual machine (VM)	A computer that uses software components, but acts like a physical machine. A virtual machine resides on a host machine.
Active/active	Two load balancers working in tandem to distribute network traffic.
Active/passive	Two load balancers with one actively working and the second in listening mode to take over if the active machine fails.
Virtual IP	An IP address that can be used by multiple endpoints. It is commonly used in failover systems and for load balancing.
Storage area network (SAN)	A dedicated, high speed network of storage devices. Usually used for file shares.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	3.1 Compare and contrast security implications of different architecture models. Architecture and infrastructure concepts Cloud Third-party vendors Network infrastructure High availability Considerations Power 3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

	Infrastructure considerations Device attribute Active vs. passive Load balancer 3.4 Explain the importance of resilience and recovery in security architecture. High availability Load balancing vs. clustering Site considerations Geographic dispersion Power Generators Uninterruptible power supply (UPS) 5.3 Explain the processes associated with third-party risk assessment and management. Vendor selection
TestOut Security Pro	4.1 Protect and Maintain Data files 4.1.2 Implement redundancy

9.4.1 Redundancy (Lesson Video)

Transcript:

Security architecture resilience refers to the design and implementation of systems and networks in a way that allows them to withstand and recover quickly from disruptions or attacks. This includes redundancy, fail-safe mechanisms, and robust incident response plans. By building resilience into the security architecture, cybersecurity teams ensure that even if a breach occurs, the impact is minimized, and normal operations can be restored quickly.

Redundancy strategies are essential to disaster recovery and business continuity planning. These strategies include continuity of operations planning, which involves developing processes and procedures to ensure critical business functions can continue during and after a disruption. Let's look at some of these strategies.
High availability (HA) clustering uses redundant systems that can automatically take over operations in case of failure, minimizing downtime. High availability (HA) is crucial in IT infrastructure, ensuring systems remain operational and accessible with minimal downtime. It involves designing and implementing hardware components, servers, networking, data centers, and physical locations for fault tolerance and redundancy. In a high-availability setup, redundant hardware components, such as power supplies, hard drives, and network interfaces, reduce the risk of failure by allowing the

system to continue functioning if one component fails. Servers are often deployed in clusters or paired configurations, which allows automatic failover from a primary server to a secondary server in case of an issue.
Power redundancy ensures critical infrastructure, such as data centers, has backup power sources to continue operations during an outage. All types of computer systems require a stable power supply to operate. Electrical events, such as voltage spikes or surges, can crash computers and network appliances, while loss of power from under-voltage events or power failures will cause equipment to fail. Power management means deploying systems to ensure that equipment is protected against these events and that network operations remain uninterrupted or recover quickly.
Vendor diversity is essential for several reasons, offering benefits not only in terms of cybersecurity but also in business resilience, innovation, and competition. Relying on a single vendor for all software and hardware solutions can create a single point of failure. The entire infrastructure may be at risk if a vulnerability is discovered in that vendor's products. Vendor diversity mitigates the risk associated with vendor lock-in. It ensures that an organization's operations aren't solely reliant on one vendor's products or services. Diverse vendors bring different perspectives, ideas, and technologies.
Vendor diversity promotes healthy competition in the market, which can lead to better pricing, improved product features, and higher-quality customer support. Different vendors offer unique solutions that cater to specific needs, and having a diverse vendor ecosystem allows organizations to choose the best fit for their requirements. Vendor diversity helps spread the risk associated with
potential product or service failures, security breaches, and other issues. In some industries, regulations or standards may require organizations to maintain vendor diversity to ensure compliance and reduce the risk of supply chain disruptions or security breaches.
Defense-in-depth is a comprehensive cybersecurity strategy that emphasizes the implementation of multiple layers of protection to safeguard an organization's information and infrastructure. This approach is based on the principle that no single security measure can completely protect against all threats. By deploying a variety of defenses at different levels, organizations can create a more resilient security posture that can withstand a wide range of attacks. For example, a defense-in-depth strategy might include perimeter security measures such as firewalls and intrusion detection systems to protect against external threats. Organizations can implement segmentation, secure access controls, and traffic monitoring at the network level to prevent unauthorized access and contain potential breaches. Endpoint security solutions, such as antivirus software and device hardening, help protect individual devices. At the same time, regular patch management ensures that software vulnerabilities are addressed promptly.
Regular testing, including tabletop exercises, failover tests, and simulations, is essential to identify vulnerabilities, evaluate response plans, and improve redundancy measures. By incorporating redundancy strategies, organizations can reduce risks, minimize downtime, and ensure the continuity of their critical business functions.
Testing high availability, load balancing, and failover technologies is critical. It assesses the ability to remain operational during heavy workloads, component failures, or scheduled maintenance. Load testing incorporates specialized software tools to validate a system's performance under expected or peak loads and identify bottlenecks or scalability issues. Failover testing focuses on validating failover processes to ensure a seamless transition between primary and secondary infrastructure. Testing monitoring systems validate effective detection and response to failures and performance issues. Robust testing practices allow organizations to ensure that high availability, load balancing, and failover technologies effectively fulfill their purpose to minimize unexpected outages and maximize performance.
That's it for this lesson. In this lesson, we discussed various redundancy strategies, including high availability, power redundancy, vendor diversity, and defense in depth. We reviewed the importance of testing high availability, load balancing, and failover technologies.

9.4.2 Redundancy Facts

This lesson covers the following topics:

- Implement secure network designs
- Manage redundant power options

Implement Secure Network Designs

The following table describes secure network designs.

Secure Network Designs	Description
Load balancing	A process that distributes processing among multiple nodes.
Active/active	Two load balancers working in tandem to distribute network traffic.
Active/passive	Two load balancers with one actively working and the second in listening mode to take over if the first one becomes unavailable.
Power scheduling	Power scheduling is used to configure an active redundancy. This sends power to networks when a power facility goes down. Power scheduling prevents total loss of power during catastrophic events.
Virtual IP (VIP)	An IP address that is not assigned to an endpoint. VIP is used for load balancing. It typically uses NAT IP address assignment.
Geographic dispersal	The use of multiple locations to store data to mitigate downtime due to location.
Multipath	A fault-tolerance technique that gives multiple physical paths between a CPU and a mass-storage appliance.

Manage Redundant Power Options

Redundant power options are vital. A network without power is useless. Common power options found in datacenters include:

Uninterrupted power supply (UPS). A UPS is a stand-alone bank of batteries that allows for the graceful shutdown of network appliances when power goes out.

Generator. A generator is a large scale device that provides power for an extended period of time. Normally between 24 and 48 hours.

Dual supply. A dual power supply is common in network appliances like servers and firewalls. It allows for one failure and hot-swapping.

Managed power distribution unit (PDU). A managed power distribution unit is a rack-mounted unit that distributes power on a large scale such as a data center.

9.4.3 Hardware Clustering (Lesson Video)

Transcript:

I'll spend a few minutes talking about clustering. Clustering can be an effective way to implement a disaster-recovery plan as well as a way to improve your productivity. A cluster is a group of interconnected servers, also known as nodes, that appear to be a single system to the operating environment. Although clustering can be done on virtual machines, I'll focus this lesson on the unique characteristics of clustering on physical computers.

Hardware clustering provides several benefits. For example, when you use clustering, the throughput and response time are dramatically improved. Since the cluster's nodes appear to be one system, if one node fails, the others in the cluster still provide the services you need and redistribute the workload among the remaining servers. To provide additional performance, more nodes can be added. Theoretically, there's no limit to the number of nodes you can have in your cluster. But you must have software that supports clustering. This software could be built into the operating system itself, such as with Windows Server 2019. In other cases, you may have to purchase a program to help you set up and manage your clusters.

With that introduction, I'll show you how clustering works.

In a typical clustering implementation, there are at least two nodes. They can be connected in multiple ways in order to act seamlessly with each other. First, since the nodes provide services to the workstations that reside on the production network, the clustered nodes are directly connected to the production network. For example, a user at this workstation can access the services on the clustered nodes through the production network.

To increase performance, clustered nodes often have a second network card, allowing them to also be connected to each other through a dedicated network. As you can see here, this network is isolated from the production network. This means that the clustered nodes connected to this dedicated network don't have to compete for bandwidth with the production traffic. When you work with high-availability clusters, this network is also referred to as a heartbeat network. I'll talk more about that in a bit. But, keep in mind that although this is the ideal setup, you don't have to use this second dedicated network. Instead, you could choose to do your clustering over the production network.

Depending on the cluster type you choose and the cluster's purpose, the nodes in your cluster could also share a common storage that's accessible through a storage area network, or SAN. When used, a SAN is often connected to the cluster using fiber-channel connections, which are fiber optic. These connections allow the devices to communicate very quickly with the shared storage. In this setup, the shared storage appears to the operating system as if it were storage installed within the server itself. But in reality, both servers share the same disk storage. This has important implications. It means that whenever Server A writes information to the SAN, it's immediately available to Server B and vice versa.

When planning your cluster, keep in mind that there are a few different types of clusters you can implement.

One commonly used cluster is called a high-availability cluster, or HA cluster. This type of cluster is also known as a failover cluster. The idea behind this specific cluster type is to eliminate downtime when a computer system in the cluster fails. Although the most common size of an HA cluster is two nodes, an HA cluster can have many more.

A high-availability cluster typically uses what's known as an active/passive configuration. With this type of configuration, the active server, or primary server, provides the services to the production network while the passive server, or standby server, waits in the background. If the primary server fails, the passive server becomes the new active server and provides the services to the production network. In this type of configuration, the passive node must be a fully redundant instance of the active node and use the same shared storage. This way, any node in the cluster has access to the same data.

To monitor when the passive server should take over, HA clusters make sure that the other servers in the cluster are alive by sending what are called heartbeats over a dedicated heartbeat network. For example, Server A, our primary server, continually lets Server B know that it's up and running by sending these heartbeats. If Server A fails, Server B no longer hears a heartbeat coming from Server A and assumes that Server A has gone down. It immediately takes over and becomes the new active server. This is possible because both servers use the same shared storage.

Depending on how close together your heartbeat intervals are set, it may take only a few seconds for a passive server to start providing the services that the active server was providing. As such, the user usually notices little to no downtime. But, keep in mind that whatever was in RAM on the failed server is lost. If the server that failed is fixed and brought back online, it becomes the passive server and listens for heartbeats from the current active server.

Another type of cluster you can use is called a load-balancing cluster, which works differently from a high-availability cluster. In a load-balancing cluster, all nodes are always active participants. This is known as an active/active configuration. In this type of cluster, all computers share in the processing workload. In a way, you can think of a load-balancing cluster as a type of supercomputer system. In other words, the processing tasks are distributed among all the nodes within the cluster. Companies that provide web server access to a large clientele typically implement load-balancing clusters to assign the many different queries to different nodes. This optimizes the responses to these requests.

Let's look at an example.

First, notice that instead of using a SAN to provide a common disk storage, each computer has its own disk storage. This isn't a requirement of a load-balancing cluster, but since a SAN can be expensive, some companies might not choose to use them. Still, using a SAN is probably the fastest and most effective way to implement a cluster when feasible.

In some cases, load-balancing clusters might also have a separate device known as a load balancer, which is used to determine which cluster node gets the current request. Load balancing uses an algorithm to determine which server in the cluster should service the request.

Similar to the example, some implementations use a round-robin approach where Server A gets the first request, Server B gets the second, and Server A gets the third.
The systems in a load-balancing cluster can be loosely linked or tightly linked. The tighter the link, the more they act as one computer system. In a loosely linked cluster, each system operates autonomously, but also in conjunction with the other systems at the same time. In a tightly linked system, the systems function as one system called a supercomputing cluster. These systems pool their CPU and storage resources, and they might even pool their memory together so that various processing tasks are distributed between the cluster members.
A key thing to remember when you implement clustering is that the more tightly integrated the systems, the more identical the hardware needs to be.
If you're using a loosely linked cluster, you can use hardware that's slightly more disparate. For example, you can use servers from different manufacturers. But, to implement a tightly linked cluster, the systems need to be identical. So, they should be from the same manufacturer, the same make and model, same processor, same amount of storage, same amount of RAM, and so on.
I don't have time to go into specific clustering implementations on various operating systems. Just be aware that most of your commonly used network operating systems do have some type of clustering solution available, whether it's built into the product itself or whether it's available from a third party.
That's it for this lesson. In this lesson, we gave you an overview of how clustering works. We talked about the role of a cluster and several clustering implementations. Then we looked at high-availability and load-balancing cluster types, and we talked about how you create a supercomputing cluster, or tightly linked cluster, from a load-balancing cluster.

9.4.4 Clustering Facts

This lesson covers the following topics:

- Virtual IP
- Application clustering

Where load balancing distributes traffic between independent processing nodes, clustering allows multiple redundant processing nodes that share data with one another to accept connections. This provides redundancy. If one of the nodes in the cluster stops working, connections can fail over to a working node. To clients, the cluster appears to be a single server. A load balancer distributes client requests across available server nodes in a farm or pool. It is generally associated with managing web traffic, whereas clusters provide redundancy and high availability for systems such as databases, file servers, etc.

Virtual IP

For example, an organization might want to provision two load balancer appliances so the other can handle client connections if one fails. Unlike load balancing with a single appliance, the public IP used to access the service is shared between the two instances in the cluster. This arrangement is called a virtual IP or shared or floating address. The instances are configured with a private connection, on which each is identified by its "real" IP address. This connection runs a redundancy protocol, such as Common Address Redundancy Protocol (CARP), enabling the active node to "own" the virtual IP and respond to connections. The redundancy protocol also implements a heartbeat mechanism to allow failover to the passive node if the active one should suffer a fault.

- With active/passive clustering, if one node is active, the other is passive. The most significant advantage of active/passive configurations is that performance is not adversely affected during failover. However, there are higher hardware and operating system costs because of the unused capacity.
- An active/active cluster means that both nodes are processing connections concurrently. This allows the administrator to use the maximum capacity from the available hardware while all nodes are functional. In the event of a failover, the workload of the failed node is immediately and transparently shifted onto the remaining node. At this time, the workload on the remaining nodes is higher, and performance is degraded.

Application Clustering

Clustering is also very commonly used to provision fault-tolerant application services. If an application server suffers a fault in the middle of a session, the session state data will be lost. Application clustering allows servers in the cluster to communicate session information to one another. For example, if a user logs in on one instance, the next session can start on another, and the new server can access the cookies or other information used to establish the login.

9.4.5 Incorporate Redundancy Strategies

9.4.6 Practice Questions (Section Quiz)

q_redundancy_dual_power_secp8

To prevent server downtime, which of the following components should be installed in a server system redundantly?

Answers:

***Uninterrupted power supply**

CD or DVD drive

RAM modules

Floppy disk drive

Explanation:

To prevent server downtime, you should install uninterrupted power supplies in a server system. If one fails, the other can immediately take over, allowing the server to remain running.

Because it isn't a critical component, a redundant CD or DVD drive probably isn't necessary. Unless data was shared from a disc in the drive, a failed CD or DVD drive probably won't affect the server's functionality.

With most motherboards, there's no way to install redundant RAM modules.

Like CD or DVD drives, the floppy disk drive isn't a critical component. A failed floppy disk drive won't bring the server down.

q_redundancy_hot_swap_secp8

You have been asked to deploy a network solution that includes an alternate location where operational recovery is provided within minutes of a disaster.

Which of the following strategies would you choose?

Answers:

Hot spare

***Hot swapping**

Cold site

Warm site

Explanation:

Hot swapping is a complete disaster recovery facility that could be fully operational within hours or minutes in the event of a disaster. This includes maintaining redundant hardware and up-to-date network data.

A warm site is a remote network location that maintains a backup of data, but it is not always current. Data may be days or weeks old, depending on backup procedures.

A cold site provides a space and sometimes hardware in an alternate location that can be configured when needed. Returning to an operational state may take days.

A hot spare is a redundant hardware component used as a failover solution.

q_redundancy_load_balance_active_secp8

A growing e-commerce company is considering various strategies to ensure its web servers can handle sudden traffic surges without an impact on site availability.

The company is assessing methods that distribute incoming traffic across multiple servers.

Which strategy should the company implement to dynamically allocate the load based on real-time traffic and server conditions?

Answers:

***Active load balancing**

Passive load balancing

Active server pages

Active Directory

Explanation:

Active load balancing constantly monitors server loads and dynamically distributes client requests across multiple servers based on their real-time conditions and capacities. This process is an effective strategy to manage sudden traffic surges and maintain high availability.

Passive load balancing distributes client requests evenly across servers without considering their current loads.

Active server pages are a server-side scripting technology that creates dynamic web pages. While it can improve the website's functionality, it does not directly impact server load distribution.

Active Directory is a service for managing network domains and resources. While Active Directory is a crucial component in many networks, it doesn't contribute to load balancing or handling traffic surges.

q_redundancy_load_balance_passive_secp8

A growing e-commerce company wants to implement a strategy that evenly distributes incoming traffic across multiple servers without constantly monitoring server loads or making adjustments based on real-time conditions.

Which strategy should this company implement to manage load distribution in this manner?

Answers:

***Passive load balancing**

Active Directory

Active Server Pages

Active load balancing

Explanation:

Passive load balancing distributes client requests evenly across servers without considering their current loads. This method aligns with the company's desire to manage load distribution without constantly monitoring real-time conditions.

Active Directory is a service for managing network domains and resources. While essential in many networks, Active Directory doesn't contribute to load balancing or managing traffic surges.

Active Server Pages is a server-side scripting technology that creates dynamic web pages. Active Server Pages can enhance a website's functionality, but does not directly impact server load distribution.

Active load balancing monitors server loads, dynamically distributing client requests based on real-time conditions. While effective for managing traffic surges, this approach does not align with the company's preference for an even distribution method without real-time adjustments.

q_redundancy_load_balance_secp8

A systems engineer is designing a new IT infrastructure for a company that provides a highly used online service. The company wants to ensure that its service communications are efficient and available around the clock.

Which feature should the engineer primarily consider during the design process?

Answers:

***Load balancing**

Multipath

Generators

Replication

Explanation:

Load balancing helps to distribute network or application traffic across several servers, preventing any single server from becoming a bottleneck. This concept is crucial for designing a high-availability system.

Multipath is a fault-tolerance technique that provides multiple physical paths between a CPU and a mass-storage appliance. However, it does not make sure that service communications are efficient and available around the clock.

Generators can be important for maintaining power supply in case of outages, but they do not make sure that service communications are efficient and available around the clock.

Replication is a method of copying data to ensure its availability in case of a hardware failure. However, it does not make sure that service communications are efficient and available around the clock.

q_redundancy_pdu_secp8

A security consultant assesses a company's server room to determine how well it can maintain operations during power interruptions. The consultant evaluates the integration of power distribution units (PDUs) and backup power generators within the security architecture.

Considering the goal of ensuring resilience and recovery in the server room during power interruptions, which primary role does the backup power generator play in conjunction with the PDU?

Answers:

It ensures power load balancing across multiple servers.

It instantly supplies power to PDUs to prevent any momentary lapse during an outage.

***It provides a prolonged source of power to the PDUs after the UPS system depletes its immediate resources.**

It filters and stabilizes the power before it is distributed by the PDU.

Explanation:

A backup power generator supplies prolonged power to PDUs once UPS systems deplete their immediate resources, ensuring continuous operation of servers and essential devices during extended power outages.

While PDUs ensure efficient power load balancing across servers, maintaining consistent energy distribution, backup power generators mainly act as a reserve energy source and are not primarily responsible for balancing.

Though backup power generators serve as a long-term power solution, they do not instantly deliver power in emergencies. Instead, UPS systems serve to provide immediate power responses, bridging the gap until generators activate.

Although it is important to have stable power, backup power generators focus on supplying continuous energy during outages. They do not primarily work to filter or stabilize the power.

q_redundancy_redundancy_secp8

What is the primary security feature that can be designed into a network's infrastructure to protect and support availability?

Answers:

Periodic backups

***Redundancy**

Fiber optic cables

Switches instead of hubs

Explanation:

Redundancy is the primary security feature that can be designed into a network's infrastructure to protect and support availability. This is because it identifies single points of failure.

Periodic backups are better than no backups, but real-time and off-site backups are better protection for availability.

Fiber optic cables are not a real protection for a network's availability, as eavesdropping protection is the only security benefit they provide.

Switches are better than hubs, but there are infrastructure security measures that provide more significant protections for availability.

q_redundancy_ups_01_secp8

A data center manager is evaluating the resilience and recovery capabilities of a company's server room. The manager wants to ensure that in the event of power fluctuations or outages, the company's servers remain operational and maintain data integrity.

The manager focuses on the role of power distribution units (PDUs) and Uninterruptible Power Supplies (UPSs) in this context.

In enhancing the resilience and recovery capabilities of the server room against power interruptions, which primary function does the UPS provide to the servers that directly support this goal?

Answers:

It distributes power to multiple servers simultaneously.

It filters power to remove noise and surges.

***It provides temporary power during an outage, allowing for a graceful shutdown or transition to backup generators.**

It monitors power usage and sends alerts for overconsumption.

Explanation:

A UPS provides temporary power during an outage, ensuring that servers can undergo a graceful shutdown without data loss or continue running until backup generators take over. This function is critical for resilience and recovery as it prevents sudden power loss, leading to data corruption.

While PDUs distribute power to multiple devices, their primary function is not to provide resilience during power interruptions.

Filtering power to remove noise and surges is essential for protecting equipment, but does not directly provide resilience during more prolonged power interruptions.

Monitoring power usage and sending alerts is more about managing power consumption and ensuring efficiency than providing resilience during power outages.

q_redundancy_ups_02_secp8

A corporation is experiencing frequent power failures in its data center, which are causing downtime and resulting in high recovery costs.

Which strategy could the corporation employ to minimize the impact of these power failures?

Answers:

***Implement a UPS system.**

Employ a hybrid cloud strategy.

Implement network segmentation.

Migrate to an SDN.

Explanation:

Uninterruptable power supply (UPS) systems act in the event of a power failure, allowing systems to keep running and/or to be safely shut down, minimizing downtime and associated costs.

Employing a hybrid cloud strategy can be beneficial for many reasons. Although some of the resources are maintainable locally, in a hybrid model, this does not solve the power failures at the data center.

Implementing network segmentation enhances security by separating the network into smaller parts, but does not solve the power failure issues in the data center.

Migrating to a software-defined network (SDN) allows for greater network control and may improve efficiency, but software depends on physical devices, which are experiencing power failures.

q_redundancy_ups_dual_power_secp8

An organization operates a large data center that supports critical business operations. Recently, the organization has struggled with frequent power interruptions leading to downtime and data loss.

To address this issue, the chief information security officer (CISO) decides to review the data center's resilience and recovery strategies, particularly emphasizing backup power.

To increase the resilience and recovery capabilities of the data center and ensure operations continue even during a power failure, which of the following options should the CISO consider? (Select two.)

Answers:

***Implement a UPS.**

***Deploy a dual power supply unit in each server.**

Purchase additional servers.

Enhance the firewall system.

Implement load balancing.

Explanation:

When the input power source fails, an Uninterruptible Power Supply (UPS) provides emergency power to a load. The UPS will ensure the continuous operation of the data center during a power outage, thus enhancing its resilience and recovery capabilities.

A dual power supply unit can ensure an uninterrupted power supply to the servers if one of the power supply units fails, thereby increasing the data center's resilience and recovery capabilities.

While purchasing additional servers may increase the data center's capacity, it does not address the power redundancy issue.

While enhancing the firewall system can help prevent certain types of cyberattacks, it does not contribute directly to resilience, recovery, or power redundancy.

Load balancing is a process that distributes processing among multiple nodes to ensure availability of data and applications. It does not contribute directly to power redundancy.

q_clustering_active_active_sec8

An IT architect of a medium-sized e-commerce business that operates 24/7 is planning to enhance the system's resilience and recovery capabilities.

As part of this project, the architect is considering a clustering solution for the servers. The architect's key objective is high availability and seamless customer experience, even in the event of unexpected server failures.

Which type of clustering setup would BEST meet the needs of this e-commerce business?

Answers:

***Active/Active Clustering.**

Active/Passive Clustering.

No clustering, just a single server.

Active/Passive Clustering with an equal number of active and passive nodes.

Explanation:

Active/Active Clustering is the most suitable for a 24/7 e-commerce business. Both nodes in this setup process the connections concurrently, maximizing the utilization of available resources.

While Active/Passive Clustering provides a degree of resilience, it does not fully utilize the available resources, as one node remains idle unless the active node fails.

A single server and no clustering are not ideal for a 24/7 e-commerce business since a single server does not provide any resilience or recovery options.

This type of setup would involve unnecessary costs for the business as it requires an equal number of passive nodes for each active node.

q_clustering_active_passive_sec8

What is a significant advantage of using an active/passive clustering configuration in a network?

Answers:

Both nodes process connections concurrently, maximizing hardware capacity.

***The performance is not adversely affected during failover.**

The active node can access the cookies or other information used to establish the login.

The public IP used to access the service is shared between the two instances in the cluster.

Explanation:

The performance is not adversely affected during failover is correct. In an active/passive configuration, the performance is not adversely affected during failover because the passive node is ready to take over the workload of the active node immediately if it fails.

In an active/passive configuration, one node is active while the other is passive, meaning they do not process connections concurrently. This is a characteristic of an active/active configuration.

The ability for a node to access cookies or other login information is a feature of application clustering, not specifically active/passive clustering.

The sharing of a public IP between two instances in a cluster is referred to as a virtual IP or shared or floating address, not a characteristic specific to active/passive clustering.

q_clustering_app_clustering_sec8

What is a primary function of application clustering in a network?

Answers:

***It allows servers in the cluster to communicate session information to one another.**

It enables the active node to "own" the virtual IP and respond to connections.

It distributes client requests across available server nodes.

It processes connections concurrently in an active/active cluster.

Explanation:

Application clustering allows servers in the cluster to communicate session information to one another. This means that if a user logs in on one instance, the next session can start on another, and the new server can access the cookies or other information used to establish the login.

The ability for the active node to "own" the virtual IP and respond to connections is a feature of the Common Address Redundancy Protocol (CARP), not a specific feature of application clustering.

The distribution of client requests across available server nodes is a function of a load balancer, not a specific feature of application clustering.

The processing of connections concurrently in an active/active cluster is a characteristic of the cluster configuration, not a specific feature of application clustering.

q_clustering_carp_sec8

What is the primary function of the Common Address Redundancy Protocol (CARP) in a network?

Answers:

To distribute client requests across available server nodes.

***To enable the active node to "own" the virtual IP and respond to connections.**

To communicate session information between servers in a cluster.

To process connections concurrently in an active/active cluster.

Explanation:

CARP enables the active node in a cluster to "own" the virtual IP and respond to connections. It also implements a heartbeat mechanism to allow failover to the passive node if the active one should suffer a fault.

The distribution of client requests across available server nodes is a function of a load balancer, not CARP.

The communication of session information between servers in a cluster is a feature of application clustering, not a function of CARP.

The processing of connections concurrently in an active/active cluster is a characteristic of the cluster configuration, not a function of CARP.

q_clustering_load_balancer_sec8

What is the primary function of a load balancer in a network?

Answers:

***To distribute client requests across available server nodes in a farm or pool.**

To provide a virtual IP or shared or floating address.

To communicate session information between servers in a cluster.

To process connections concurrently in an active/active cluster.

Explanation:

The primary function of a load balancer is to distribute client requests across available server nodes in a farm or pool, managing web traffic and ensuring that no single server becomes overwhelmed.

While a load balancer can be part of a setup that uses a virtual IP or shared or floating address, this is not its primary function. This is more related to the concept of clustering.

The communication of session information between servers in a cluster is a feature of application clustering, not a function of a load balancer.

The processing of connections concurrently in an active/active cluster is a characteristic of the cluster configuration, not a function of a load balancer.

9.5 Backup and Restore

As you study this section, answer the following questions:

- Why are there different backup types?
- How often should you run a full backup?
- How do incremental and differential backups differ?
- How can you implement the 3-2-1 rule?
- What are the differences between network attached storage (NAS) and storage attached network (SAN)?

In this section, you will learn to:

- Configure network attached storage.
- Implement file backups.
- Back up files with file history.
- Recover a file from file history.
- Backup a domain controller.
- Restore server data from a backup.

Key terms for this section include the following:

Term	Definition
Full backup	A back up that captures all of the data on a machine. A full backup is always the first backup you should run.
Incremental backup	A backup that contains all changes since the last incremental backup.
Differential backup	A backup that contains all changes since the last full backup.

Snapshot	An instant copy of an individual computer. Snapshots are normally used on virtual machines (VMs) when changes may need to be reverted.
NAS	A network storage appliance often used to store backups or other files.
SAN	A network of fast storage appliances. A SAN stores file shares and other data that needs to be accessed quickly.
Offsite storage	A location where files are stored that is away from the physical office space where the data is created. Offsite storage is part of 3-2-1 rule.
Scalability	The ability to increase or decrease data storage space.
Restoration order	Pre-planned order in which servers will be restored following a disastrous event. The order is determined by the server's importance to the company's operation.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	3.4 Explain the importance of resilience and recovery in security architecture. Backups Onsite/offsite Frequency Encryption Snapshots Recovery Replication Journaling
TestOut Security Pro	4.0 Data Security 4.1 Protect and Maintain Data files 4.1.1 Perform data backups and recovery

9.5.1 Backup Types (Lesson Video)

Transcript:

In this lesson, we'll discuss backing up data.

Backing up your data is absolutely critical. It must be done consistently, and strategically. It's also essential to verify backups to ensure that they actually work and you can restore your data.
You need to know three different types of backups. The first is the full backup, the second is the incremental backup, and the last is the differential backup. But before we dive into these, let's quickly talk about a key component used with backups called the archive bit.
The archive bit is a file attribute that is either on or off. It's used by backup systems to determine if the file has been archivedâ€""that is, has been backed up. Also, archive bit identifies if a file has been modified since the last backup. This is important when used with certain types of backups.
With that, let's explore the different types of backups.
Let's first discuss a full backup. A full backup backs up everything. It doesn't matter whether the archive bit is set or not on any file. However, as it backs up each file, it clears the archive bit to indicate that the file is backing up.
In this way, you can determine if a file has been modified at a later point in time. If you were to modify a file that has been backed up, the archive bit would be reset, indicating to the backup software that this file has been modified since the last time it was backed up.
If a file is not modified after the full backup, the archive bit is left in a cleared state. This is important because other types of backups, such as incremental, take into consideration if this archive bit has been cleared when deciding whether it should back up a file or not.
An incremental backup backs up everything since the last full backup or since the last incremental backup. To determine whether or not a file has been modified since the last full or incremental backup, it looks at the archive bit. If the archive bit is turned on, that tells the backup software that the file has been modified since the last full or incremental backup, and needs to be backed up again. Once it's backed up the file, the archive bit is cleared.
Let's run through an example. Suppose you perform a full backup of the entire file system on Monday. As a result, every file in the file system has been backed up, and the archive bit on each file is cleared. Then, on Tuesday, you perform an incremental backup. Because the incremental backup only backs up files with the archived bit set, it will only backup files that have changed since Monday. Accordingly, the incremental backup you made on Tuesday will be relatively small because it will only contain one day's worth of changes. After backing up the changed files, the incremental backup will clear the archive bit again on the files that were backed up.
If you perform another incremental backup on Wednesday, it will once again look for any files with enabled archive bits, indicating that they've been modified since either the last full backup on Monday or since the last incremental backup on Tuesday. Therefore, Wednesday's incremental backup also only contains one day's worth of changes. Likewise, if you run another incremental backup on Thursday, then only the changes that have been made since Wednesday will be backed up. The same is true on Friday.
The advantage of this strategy is that the daily incremental backups finish relatively quickly because we're only backing up one day's worth of changes. The full backup still takes some time to complete because all files are being backed up. For this reason, the full backup is usually executed over the weekend when the system is not heavily used. The incremental backups can occur each weekday. Because they don't take much time to complete, they typically don't interfere with day-to-day work.
Incremental backups have one significant drawback. The incremental backup strategy is the slowest type of backup for restoring data. For example, suppose we ran a full backup on Monday and then ran an incremental on Tuesday, Wednesday, Thursday, and Friday. Then the server crashes on Saturday, and we need to restore all the data back to the server after we've recovered from the crash.
The first thing that you must do is restore the first full backup from Monday. Then you must restore every incremental backup in the order they were created. First, you must restore Monday's incremental backup, followed by Tuesday's backup, then Wednesday, Thursday, and Friday's backups. Each backup must be restored in the proper sequence to return the system to where it was before the crash. This can take hours or even days to complete, depending how much data is involved.
Another option is to use differential backups instead of incremental. A differential backup looks for files that have been modified since the last full backup. It does this by evaluating if the archive bit has been set on the files. If the archive bit

is clear, it will not back up the file because it assumes that it has not been changed since the last full backup. If the archive bit is set, the differential backup assumes the file has been modified and needs to be backed up.
The key difference is that a differential backup does not clear the archive bit after backing up a file. Therefore, it backs up everything modified since the last full backup, but not since the last differential backup. This has advantages and disadvantages.
Suppose you perform a full backup on Monday just as with the incremental backup strategy. This clears all archive bits on all files. Then, we perform a differential backup on Tuesday and Wednesday. Because the differential backup on Tuesday does not clear any archive bits, Wednesday's differential backs up all the files that have been modified since Monday. Thursday's differential backup includes all the changes made on Monday, Tuesday, and Wednesday. Friday's differential backup includes all the changes made on Monday, Tuesday, Wednesday, and Thursday.
As you progress through the backup schedule, each differential backup takes progressively longer because it's backing up more data each time. The first differential backup completes in the same amount of time as an incremental would because we're only backing up one day's worth of data. However, each subsequent differential takes longer.
The advantage of differential backups is the speed of restoring data. With differentials, all we do is restore the full backup followed by the last differential backup. For example, if a server were to crash on Saturday, we would restore the full backup from Monday and then the last differential backup from Friday. That's all we do because the last differential backup contains a backup of every single file that has been changed since the last full backup.
You should never mix differential and incremental backups together with a full backup. If you mix full backups with incremental and differential backups, you're going to have problems because of the way these different types of backups handle the archive bit. As a result, files that should have been backed up will be missed.
Another alternative for backing up data is to create a system image. Imaging programs create a bit-level mirror of a particular hard disk or partition.
Newer versions of Windows, and all versions of Linux, include this functionality. Typically, a system image is created on a defined schedule. Because a system image backs up the entire hard disk or partition, they are rather slow to create. However, they are very fast to restore.
All you do is restore the image to the same hardware, or even to a new piece of hardware if you have a major hardware malfunction, and the system is back and available again in exactly the state it was when the image was created. This can be very useful in the event of a malware infection. Most of the time it's easier to re-image a machine that's been infected with malware rather than try to get rid of the malware.
There are other types of data to back up. For example, if you're managing a domain controller in a Windows network, then you need to backup not only the files on the server but all your active directory information as well. To do this, you can create a special type of backup on the domain controller called a system state backup.
If you are backing up sensitive information, such as your security logs, then you might want to consider the location of the backup. Instead of backing that data up to a device that can be modified in some way, find a more secure option. For example, it could be problematic if you were to back up sensitive information to a network share or perhaps a USB flash drive because of the potential of tampering. Instead, you should back up confidential information to a write-once type of media, such as a recordable DVD.
A backup is worthless unless you can successfully restore data from it. You should run periodic tests and restore from your backups in a lab environment, just to make sure that everything works should a crisis occur.
That's it for this lesson. In this lesson, we talked about backup and restore. We first discussed the importance of backing up your data. We then examined the three backup strategies: full, incremental, and differential. We then mentioned using a system image to backup data. We talked about backing up the system state on the domain controller. We discussed backing up sensitive information to write-once medium. Then we emphasized the importance of verifying backups to make sure that they actually work.

9.5.2 Backup Storage Options (Lesson Video)

Transcript:

Backups play an essential role in asset protection by ensuring the availability and integrity of an organization's critical data and systems. By creating copies of important information and storing them securely in separate locations, backups are a safety net in case of hardware failure, data corruption, or cyberattacks such as ransomware.

Many dynamics influence data backup frequency requirements, including data volatility, regulatory requirements, system performance, architecture capabilities, and operational needs. Organizations with highly dynamic data or stringent

regulatory mandates may opt for more frequent backups to minimize the risk of data loss and ensure compliance. Conversely, businesses with relatively stable data or less rigorous regulatory oversight might choose less regular backups, balancing data protection, data backup costs, and maintenance overhead. Ultimately, the optimal backup frequency is determined by carefully assessing an organization's regulatory requirements, unique needs, risk tolerance, and resources.

The need for on-site and off-site backups must be balanced, as they're crucial in securing critical data and ensuring business continuity. On-site backups involve storing data locally—in the same location as the protected systems—on devices such as hard drives or tapes to provide rapid access and recovery in case of data loss, corruption, or system failures. On the other hand, off-site backups involve transferring data to a remote location to ensure protection against natural disasters, theft, and other physical threats to local infrastructure, as well as catastrophic system loss resulting from ransomware infection.

Ransomware poses a significant threat to businesses and organizations by encrypting vital data and demanding a ransom for its release. In many cases, ransomware attacks also target backup infrastructure, hindering recovery efforts and further exacerbating the attack's impact. Perpetrators often employ advanced techniques to infiltrate and compromise both primary and backup systems, rendering them useless when needed. Organizations can implement several strategies to defend against this risk, such as maintaining air-gapped backups physically disconnected from the network, thereby actively preventing ransomware from accessing and encrypting them.

Critical recovery validation techniques play a vital role in ensuring the effectiveness and reliability of backup strategies. Organizations can identify potential issues and weaknesses in their data recovery processes by testing backups and making necessary improvements.

One common technique is the full recovery test, which involves restoring an entire system from a backup to a separate environment and verifying the fully functional recovered system. This method helps ensure that all critical components, such as operating systems, applications, and data, can be restored and function as expected.

Another approach is the partial recovery test, where selected files, folders, or databases are restored to validate the integrity and consistency of specific data subsets. Organizations can perform regular backup audits, checking the backup logs, schedules, and configurations to ensure backups are created and maintained as intended and required.

Furthermore, simulating disaster recovery scenarios, such as hardware failures or ransomware attacks, provides valuable insights into an organization's preparedness and resilience. Recovery validation strategies are essential because backups can be completed with "100% success" but mask issues until the backup set is used for recovery.

Encryption of backups is essential for various reasons, primarily data security, privacy, and compliance. By encrypting backups, organizations add an extra layer of protection against unauthorized access or theft, ensuring that sensitive data remains unreadable without the appropriate decryption key. This is particularly crucial for businesses dealing with sensitive customer data, intellectual property, or trade secrets, as unauthorized access can lead to severe reputational damage, financial loss, or legal consequences.

Copies of sensitive data stored in backup sets are often overlooked, so many industries and jurisdictions have regulations that mandate the protection of sensitive data stored in backups. Encrypting backups helps organizations meet these regulatory requirements and avoid fines, penalties, or legal actions resulting from noncompliance.

That's it for this lesson. In this lesson, we talked about data backups and backup frequency. We discussed the value of maintaining on-site and off-site backups. We talked about identifying potential issues and weaknesses by performing recovery validation. We also discussed the importance of encrypting backups.

9.5.3 Backup Types and Storage Facts

This lesson covers the following topics:

- Enterprise backups
- Data deduplication
- Backup frequency
- Snapshots
- Replication and journaling

Enterprise Backups

In an enterprise setting, simple backup techniques often prove insufficient to address large organizations' unique challenges and requirements. Scalability becomes critical when vast amounts of data must be managed efficiently. Simple backup methods may struggle to accommodate growth in data size and complexity.

Performance issues caused by simple backup techniques can disrupt business operations because they slow down applications while running and typically have lengthy recovery times. Additionally, enterprises demand greater granularity and customization to target specific applications, databases, or data subsets, which simple techniques often fail to provide.

Compliance and security requirements necessitate advanced features such as data encryption, access control, and audit trails that simplistic approaches typically lack. Moreover, robust disaster recovery plans and centralized management are essential for an enterprise backup strategy. Simple backup techniques might not support advanced features like off-site replication, automated failover, or streamlined management of the diverse systems and geographic locations that comprise a modern organization's information technology environment.

Critical capabilities for enterprise backup solutions typically include the following features:

- Support for various environments (virtual, physical, and cloud)
- Data deduplication and compression to optimize storage space
- Instant recovery and replication for quick failover
- Ransomware protection and encryption for data security
- Granular restore options for individual files, folders, or applications
- Reporting, monitoring, and alerting tools for effective management
- Integration with popular virtualization platforms, cloud providers, and storage systems

Data Deduplication

Data deduplication describes a data compression technique that optimizes storage space by identifying and eliminating redundant data. It works by analyzing data blocks within a dataset and comparing them to find identical blocks. Instead of storing multiple copies of the same data, deduplication stores a single copy. It creates references or pointers to that copy for all other instances. Deduplication can be performed at different levels, such as file-level, block-level, or byte-level. Deduplication significantly minimizes storage requirements and improves data transfer efficiency, particularly in backup and data replication processes, by reducing the amount of duplicate data stored.

Backup Frequency

Many dynamics influence data backup frequency requirements, including data volatility, regulatory requirements, system performance, architecture capabilities, and operational needs. Organizations with highly dynamic data or stringent regulatory mandates may opt for more frequent backups to minimize the risk of data loss and ensure compliance. Conversely, businesses with relatively stable data or less stringent regulatory oversight might choose less frequent backups, balancing data protection, data backup costs, and maintenance overhead. Ultimately, the optimal backup frequency is determined by carefully assessing an organization's regulatory requirements, unique needs, risk tolerance, and resources.

Snapshots

Snapshots play a vital role in data protection and recovery, capturing the state of a system at a specific point in time. Virtual Machine (VM), filesystem, and Storage Area Network (SAN) snapshots are three different types, each targeting a particular level of the storage hierarchy.

VM snapshots, such as those created in VMware vSphere or Microsoft Hyper-V, capture the state of a virtual machine, including its memory, storage, and configuration settings. This allows administrators to roll back the VM to a previous state in case of failures, data corruption, or during software testing.

Filesystem snapshots, like those provided by ZFS or Btrfs, capture the state of a file system at a given moment, enabling users to recover accidentally deleted files or restore previous versions of files in case of data corruption.

SAN snapshots are taken at the block-level storage layer within a storage area network. Examples include snapshots in NetApp or Dell EMC storage systems, which capture the state of the entire storage volume, allowing for rapid recovery of large datasets and applications.

Replication and Journaling

Replication and journaling are data protection methods that ensure data availability and integrity by maintaining multiple copies and tracking changes to data. Replication involves creating and maintaining exact copies of data on different storage systems or locations. Organizations can safeguard against data loss due to hardware failures, human errors, or malicious attacks by having redundant copies of the data. In the event of a failure, the replicated data can be utilized to restore the system to its original state.

A practical example of replication is database mirroring, where an organization maintains primary and secondary mirrored databases. Any changes made to the primary database are automatically replicated to the secondary database, ensuring data consistency and availability if the primary database encounters any issues.

On the other hand, journaling records changes to data in a separate, dedicated log known as a journal. Organizations can track and monitor data modifications and revert to previous states if necessary. Journaling is beneficial for data recovery in system crashes. After completing the full system backup, it enables the system to identify and undo any incomplete transactions that might have caused inconsistencies or replay transactions. This provides greater granularity for restores and greatly minimizes data loss. A practical example of journaling is using file system journaling, such as the Journaled File System (JFS) or the New Technology File System (NTFS), with journaling enabled. These file systems record all changes made to files, allowing for data recovery and consistency checks after unexpected system shutdowns or crashes.

Remote journaling, SAN replication, and VM replication are advanced data protection methods that maintain data availability and integrity across multiple locations and systems. Remote journaling creates and maintains a journal of data changes at a separate, remote location, allowing for data recovery and ensuring business continuity in case of local failures, natural disasters, or malicious attacks.

SAN replication duplicates data from one SAN to another in or near real-time, providing redundancy and protection against hardware failures, human errors, or data corruption. This technique involves synchronous replication, which guarantees data consistency, and asynchronous replication, which is more cost-effective but slightly less stringent in consistency.

Meanwhile, VM replication creates and maintains an up-to-date copy of a virtual machine on a separate host or location, ensuring that a secondary VM can quickly take over the workload in the event of a primary VM failure or corruption. By implementing these methods, organizations can bolster their data protection strategies, safeguarding against various risks and ensuring the availability and integrity of their critical data and systems.

9.5.4 Configure Network Attached Storage (Demo Video)

Transcript:

Network attached storage, or NAS, is a file-level storage device on your network. These devices can be useful when you share data or even back it up. Generally, there'll be several drives in the NAS device for redundancy. But if you want to make sure data isn't lost due to drive failure, a cloud sync is suggested.

TrueNAS, which is also known as free NAS, is open-source software. You can install this software on a physical server or a virtual machine.
TrueNAS makes it very easy to set up, as all we've done up to this point is install it, set an IP, and set a root password. In order to use any disks for storage, we must configure a storage pool. To do so, go to 'Storage' and then 'Pools'. Then click 'ADD' and 'CREATE POOL'. Next, we're going to give it a name and select the two available disks we have. Most likely, a production NAS will have more storage and more available disks then what we're showing. Since there are only two disks, we're going to set up a 'Mirror' vs. Stripe. Click 'CREATE', 'Confirm', and 'CREATE POOL'.
Now that our pool is set up, we must create a Windows share. Go to 'Sharing' and then 'Windows Shares (SMB)'. The drop-down menu at the top of the main folder will be selected since that was the name set on the storage pool. Nothing else is needed, but you could add a description if you'd like for documentation. Click 'SUBMIT'. A box will appear asking you to enable the SMB service since it wasn't on before.
Next, a user needs to be set up so some form of authentication can happen to access this newly created share. We'll set up a basic user under 'Accounts' and then 'Users'. Click 'ADD'. Our name here will be 'backup'. The same name will be used for the username as well. Enter a password and confirm it. As we scroll down, we need to give this user permission to our new share directory. Drop down the folder and select 'main'. All of the home directory permissions are fine for now, so we'll leave them. One thing is for certain - this user doesn't require shell login, so we want to disable this for security purposes. If you click over here and select 'nologin', it'll take care of this problem. Click 'SUBMIT' when finished. Our configuration is complete on our NAS, so were going to hop over to our Windows machine to configure the rest.
In order to have this NAS easily accessible, we're going to map it as a network drive. To do so, go to 'File Explorer'. As we expand this, we're going to go to 'This PC'. Next, the 'Computer' tab at the top will give us the option to map a network drive. Select a drive letter of your choosing. The path in the folder will be the IP address or DNS name for the NAS. We're going to use the IP as we don't have DNS set up for this right now. Type '\\192.168.30.106', as this was the IP of our NAS. Now click 'Browse'. It'll take a second, and you might have to click on the device for the username and password prompt to come up. Enter the username and password we set up for the backup user. The key thing is to check 'Remember my credentials', otherwise this mount will be asking for the password every time your computer reboots. Click 'OK' and then select the 'backup' folder. If you notice, TrueNAS creates a folder with the same username as the user, unless changed. Once we click 'Finish', our new mount will be mounted as the N: drive. Don't worry about the rest of the files in this directory, as this pertains to the username associated with the TrueNAS side. Since we may have multiple computers backing up their data here, we'll just name a folder 'computer_1'.
Now we can go set up some backup jobs. You do have the ability to use other backup solutions. However, we're just going to use the built-in File History. When we type 'backup' in the search bar here, it brings us to our backup settings. Notice when we click on 'Add a drive', it says no usable drives found. This doesn't mean it's broken. Typically, this menu looks for secondary or removable drives that are local. If we select 'More options' and 'See advanced settings', we get the option to map file history to a network location. We can then navigate to our mapped drive by going to 'This PC', clicking on the 'N:' drive, and going to the 'computer_1' folder. Click 'OK' and then after that we should be able to turn on our 'File History' now.
That's it for this demo. In this demo, we showed you how to configure a NAS, set up an SMB share, and configure a mapped network drive for File History backups.

9.5.5 Implementing File Backups (Demo Video)

Transcript:

In this demonstration, we're going to look at two ways to back up files on a Windows system—using File History and using the legacy Backup and Restore utility.

Let's take a look at setting up File History first. File History saves copies of your files so they can be recovered later if needed. Using the Windows search button, I'll look for file history and click to start it. As you can see, it's currently turned

off. There's only one button to choose from. Once it's turned on, you can select Run now to start an immediate backup of the listed locations. As you can see, the H: drive is currently selected as the storage device.
To the left, additional options are available. In addition to the Restore personal files option, which is used for retrieving copies after they've been backed up, there's Select drive. When clicked, the available devices are listed. In this case, only the H: drive is available. There's also an option to add a network location.
Another important setting is Exclude folders. Currently, the Pictures library is excluded from copies being kept. The Add and Remove buttons can be used to change this list. In Advanced settings, we can change how often copies are made and how many versions to keep. The defaults are Every hour and Forever, respectively.
The legacy Backup and Restore tool is designed to hold copies of data or an entire system as a disk image. It's recommended that the device holding the backups be external, or at least on a different storage device than the C: drive. In this case, I have a second device prepared. I'll browse to the Backup and Restore utility found under Control Panel. The tool's full name is Backup and Restore (Windows 7) since this is a legacy tool from the Windows 7 days. Once launched, I'll select Set up backup and then select the appropriate device. The H: drive is the device I intend to perform backups to, so I don't need to make any changes. I could, if necessary, click on the Save on a network button and browse to a shared location for storage. I'll leave it as the H: drive and click Next.
At this point, the system wants me to choose which files and folders to back up. Since the default settings are good for our purposes in this demo, I'll leave it on Let Windows choose and click Next. On the next screen, we see a review of the options we've selected and a new Change schedule link. Currently, the backup is run on Sunday at 7:00 PM, which is good for our needs. To finish the configuration, we click Save settings and run backup.
That's it for this demonstration on backups. We covered two ways to back up files. First, we looked at backing up using File History, and then we looked at the legacy Windows 7 Backup and Restore utility.

9.5.6 Back Up Files with File History (Simulation)

Scenario

You have recently installed a new Windows 10 computer. To protect valuable data, you need to implement file history backups on this computer.

In this lab, your task is to configure automatic backups for the Exec computer as follows:

- Save the backup to the **Backup (E:)** volume.
- Back up files **daily** .
- Keep backup files for **six months** .
- Back up the entire **Data (D:)** volume.
- Make a backup now.

Explanation

Complete this lab as follows:

1. Access the File History Backup options.
   a. Right-click **Start** and then select **Settings** .
   b. Select **Update & Security** .
   c. From the left pane, select **Backup** .

Configure and run a file history backup plan.

From the right pane, select **Add a drive** .

Select **Backup (E:)** .

Under *Automatically back up my files* , move the switch to **On** .

Select **More options** .

Under *Back up my files* , use the drop-down menu to select **Daily** .

Under *Keep my backups* , use the drop-down menu to select **6 months** .

Under *Back up these folders* , select **Add a folder** .

Double-click the **Data (D:)** volume and then select **Choose this folder** .

Select **Back up now** .

Wait for the backup to complete.

9.5.7 Recover a File from File History (Simulation)

Scenario

Susan produces your organization's monthly magazine. While working on an upcoming issue, Susan accidentally deleted significant portions of the layout image. She also made extensive changes to the cover artwork but has now been asked to discard the changes and use the original artwork.

Susan has asked you to help her recover older versions of her files in the Pictures library so she can still meet her publishing deadline.

In this lab, your task is to complete the following:

Using the Settings app, access the program needed to restore files from a current backup.

From the File History dialog, restore the following files:

File	File Version to Restore
Pictures\Layouts\June2023_Issue.jpg	Thursday, March 16, 2023 11:15 AM
Pictures\Images\coverart.jpg	Thursday, March 16, 2023 12:15 PM

Explanation

Complete this lab as follows:

Access the File History options using the Settings app.

Right-click **Start** and then select **Settings** .

Select **Update & Security** .

From the left pane, select **Backup** .

Make sure *Automatically back up my files* is set to **On** .

Select **More options** .

Scroll to the bottom of the *Backup options* dialog and select **Restore files from a current backup** .

Maximize the window for better viewing.

Restore the June2023_Issue.jpg file.

From the bottom of the File History dialog, select the **Previous version** button (left arrow) to navigate to the backups captured on **Thursday, March 16, 2023 11:15 AM** .

Double-click **Pictures** .

Double-click **Layouts** .

Select the **June2023_Issue.jpg** file.

Select the green **Restore to original location** arrow located at the bottom center.

Select **Replace the file in the destination** .
The Layouts folder where the file was restored is opened.

From the Layouts folder, right-click the **June2023_Issue.jpg** file and then select **Properties** .

Verify that the file is **115.44 MB** in size and was last modified on **March 16, 2023 at 11:15:12 AM** .

Select **OK** .

Close the Layouts window.

Restore the Coverart.jpg file.

From the top left of the *File History* dialog, select the **up arrow** to navigate to the **Home\Pictures** folder.

Select the **Previous version** button at the bottom to navigate to the backups captured on **Thursday, March 16, 2023 12:15 PM** .

Double-click **Images** .

Select the **coverart.jpg** file.

Select the green **Restore to original location** arrow located at the bottom center.

Select **Replace the file in the destination** .

Right-click the **coverart.jpg** file and select **Properties** .

Verify that the file is **1.09 MB** in size and was last modified on **March 16, 2023 at 12:15:12 PM.**

Select **OK** .

9.5.8 Backup a Domain Controller (Demo Video)

Transcript:

In this demo we are going to do a backup of our domain controller on a Windows Server. To start, you need to install the Windows backup feature. The first thing we need to do is go to Manage, click Add Roles and Features.

Now we're going to just select Default, Next, Next, past everything until we get to features. Then we're going to scroll down and go to Windows Server Backup. Select it, click next. On the confirm window I can click Install. Then that's going to go ahead and install and that may take a few minutes.
Once that's installed, we can go to our tools all the way down, scroll down to Windows Server Backup. Using Windows Server Backup you can schedule regular server backups, or you can schedule a single backup.
In this case, we're going to just choose Backup Once. However, a good disaster recovery program would do regular backups. We'll click backup once since were not scheduling a backup. Since we have not created any scheduled backups before, you must click different options.
Let's go ahead and click Custom. Click Next, we're going to click add items and we're going to select the System State. Bare Metal Recovery, will backup all the critical volumes on the computer. It will backup the operating system and all the data volumes if you are doing a bare metal recovery, also sometimes known as an all critical backup, the backup can't be saved on any of the volumes that are being backed up. As well with the system state backup, we can't save the backup on the C drive because that's where the system state exists.
So the system state includes the registry key files, Active Directory and the SYSVOL. We need to have some other place to store this backup. I've connected an external drive to this system currently, so I'm going to click OK.
Click next and you actually also have the ability to save things to remote share folders. If you have some shared storage on your network, you could actually go save your backup there.
We're going to select a local drive since it's a drive I have plugged into this local machine. It's my F drive labeled Backup. We're going to click Next, and Backup. Okay, so once this backup finishes, which it will take a while, we'll have a good backup of our active directory and the SYSVOL.
If we need to restore, we can come back to this backup later to restore the main domain controller.
In this demonstration we installed Windows server backup on your Windows Server and then backed up our domain controller.

9.5.9 Restoring Server Data from Backup (Demo Video)

Transcript:

In this demonstration, we'll explore the process of recovering user data on a server.

In our File Explorer window, let's navigate to the Desktop, where we can find the Shares folder that we previously created.

Now, let's delete this folder, making it disappear. However, imagine that later on, we realize we need it back. To retrieve it, we'll have to restore it from a backup. To initiate this process, let's go to Tools > Windows Server Backup.
In the Messages Box, we can see the available backup options. Let's click on the one we want to use and review its information. Once we're satisfied, click OK. Next, we'll click on "Recover" and then proceed with "Next."
Here, we can verify our selection to ensure it's the correct backup. Since it's the only one available, let's proceed by clicking "Next."
Note that we have several recovery options. We can recover files or folders, and if this were a Hyper-V Server, we could recover an entire virtual machine if needed. It's also possible to recover entire volumes.
Now, let's talk about applications recovery. If you have backed up a server with installed applications like Exchange or SharePoint, you can restore the entire application without having to reinstall it from scratch. However, you only need to recover the data in this case. This can be useful for transferring an application from one server to another using the backup.
The last option is to recover specific files and folders. Let's click "Next."
So, we'll navigate to the specific folder we want to recover.
After selecting it, click "Next," and then choose to restore it to the original location. Although you have the option to save it elsewhere, we want it back in its original place.
Now, let's create copies of both versions just as a precautionary measure. There are several reasons not to overwrite existing versions.
If you're restoring a file that someone claims is missing, and they provide the wrong filename, you could accidentally overwrite a file they intended to keep. This could lead to future problems.
It's a wise practice to create copies to avoid data loss or accidental overwrites. In such cases, you'd have to go back and recover even more data. Let's proceed by selecting "Next" and then "Recover."
Now, it confirms that the process is complete. Let's close and minimize the windows. You can now see that the folder has been successfully restored to its original location.
In this demonstration, we walked you through the process of restoring data using Windows Server 2022 Backup, illustrating how to utilize Windows Restore Data effectively.

9.5.10 Backup a Domain Controller (Simulation)

Scenario

You are the IT administrator for a small corporate network. You need to back up the system state of your domain controllers so that, in the event of a disaster, Active Directory is backed up. You want to configure regular backups on CorpDC4.

In this lab, your task is to perform the following using Windows Server Backup on CorpDC4:

Create a regular backup schedule for the CorpDC4 server using the following settings:

Backup items: **System State**

Backup schedule: **once per day at 1:00 a.m.**

Backup location: **\\CorpFiles\Backup**

Take an immediate backup using the following settings:

Backup items: **System State and C: drive**

Backup location: **\\CorpFiles\Backup**

Explanation

Complete this lab as follows:

Access Windows Server Backup on the CorpDC4 server.

From Hyper-V Manager, select **CORPSERVER2** .

From the Virtual Machines pane, double-click **CorpDC4** .

From the Server Manager menu bar, select **Tools** > **Windows Server Backup** .

Maximize the window for better viewing.

Create a backup schedule.

From the left pane, select **Local Backup** .

From the far right pane, under Actions, select **Backup Schedule** .

Select **Next** in the wizard.

From the Select Backup Configuration window, select **Custom** , and then select **Next** .

Select **Add items** .

Select **System state** , and then select **OK** .

Select **Next** .

Make sure **Once a day** is selected.

Using the *Select time of day* drop-down list, select **1:00 AM** , and then select **Next** .

Select **Back up to a shared network folder** , and then select **Next** .

Read the warning message, and then select **OK** .

In the Location field, enter **\\CorpFiles\Backup** and select **Next** .

Select **Finish** .

Select **Close** .

Perform an immediate backup.

From the far right pane, under Actions, select **Backup Once** .

From the Backup Options window, select **Different options** , and then select **Next** .

From the Select Backup Configuration window, select **Custom** , and then select **Next** .

Select **Add items** .

Select **System state** .

Select **Local Disk (C:)** .

Select **OK** .

Select **Next** .

Select **Remote shared folder** , and then select **Next** .

In the Location field, enter **\\CorpFiles\Backup** and select **Next** .

Select **Backup** to start the backup.

When the backup is complete, select **Close** .

9.5.11 Practice Questions (Section Quiz)

q_bkp_stor_backup-frequency_secp8

You are the IT manager of a financial services firm. Your company deals with highly dynamic data and stringent regulatory mandates.

You are tasked with determining the optimal backup frequency for your company's data.

Which of the following factors should be your top priority when deciding on the backup frequency?

Answers:

***The volatility of the company's data.**

The company's office hours.

The usability of the backup software's interface.

The brand popularity of the backup solution.

Explanation:

The volatility of the company's data is a crucial factor to consider when deciding on the backup frequency. Companies with highly dynamic data may need more frequent backups to minimize the risk of data loss.

The company's office hours do not directly impact the backup frequency. While backups are ideally scheduled during non-peak hours to minimize disruption, the frequency of backups should primarily be determined by the volatility of the data and regulatory requirements.

The usability of the backup software's interface does not impact the backup frequency and should not be a factor in this decision.

The brand popularity of the backup solution does not impact the backup frequency and should not be a factor in this decision. The focus should be on the specific needs of the company and how well the backup solution can meet these needs.

q_bkp_stor_data-deduplication_01_secp8

An enterprise seeks to optimize its backup storage space due to the increasing amounts of data it handles daily. The company wants to ensure it is not storing redundant copies of the same data, which consumes valuable storage resources.

Which technique should the company implement to solve this issue?

Answers:

***Data deduplication**

Data journaling

Data encryption

Data replication

Explanation:

Data deduplication is a data compression technique that reduces storage requirements by eliminating redundant copies of the same data. Instead of storing multiple copies of the same data, deduplication stores a single copy and creates references to that copy for all other instances.

Data journaling is a method used to keep track of changes made to the data. Data journaling is beneficial for data recovery, but does not help optimize storage space by eliminating redundant data.

While data encryption is critical for data security and compliance, it does not directly contribute to optimizing storage space.

Data replication involves creating and maintaining exact copies of data on different storage systems or locations, which can consume more storage space.

q_bkp_stor_data-deduplication_02_secp8

Which of the following statements about data deduplication is correct?

Answers:

Data deduplication increases the amount of storage space required.

Data deduplication can only be performed at the file-level.

Data deduplication reduces data transfer efficiency in backup and data replication processes.

***Data deduplication is a data compression technique that eliminates redundant data and optimizes storage space.**

Explanation:

Data deduplication is a data compression technique that works by identifying and eliminating redundant data. Instead of storing multiple copies of the same data, deduplication stores a single copy and creates references to that copy for all other instances. This process optimizes storage space and improves data transfer efficiency.

Data deduplication does not increase the amount of storage space required. On the contrary, it reduces the amount of storage space required by eliminating redundant data.

Data deduplication can be performed at different levels depending on the specific needs and capabilities of the system. This includes file-level deduplication, block-level deduplication, and byte-level deduplication.

By reducing the amount of redundant data stored, data deduplication can significantly improve the efficiency of data transfer in backup and data replication processes. This can result in faster backups and less network bandwidth usage.

q_bkp_stor_enterprise-backup_01_secp8

You are the IT manager of a large corporation. Your company has been experiencing rapid growth, and the current backup techniques are proving insufficient to address the unique challenges and requirements of the organization.

You are tasked with implementing a new enterprise backup solution.

Which of the following would be the MOST effective approach?

Answers:

Implement a simple backup method that backs up all data once per week.

***Implement an enterprise backup solution that supports various environments and offers data deduplication and compression.**

Implement a backup solution that only focuses on the most critical data and ignores less important data.

Implement a backup solution that backs up all data once per month to minimize disruption to business operations.

Explanation:

An enterprise backup solution that supports various environments and offers data deduplication and compression would address the unique challenges and requirements of a large organization. It would provide scalability, minimize disruption to business operations, meet compliance and security requirements, and offer robust disaster recovery plans and centralized management.

While implementing a simple backup method that backs up all data once per week may be simpler and less costly, it may not be sufficient for a large organization with vast amounts of data and unique challenges. It may also result in lengthy recovery times and lack the advanced features necessary for compliance and security requirements.

While it is important to prioritize critical data, ignoring less-important data could still result in significant data loss and potential operational and compliance issues. An effective backup solution should aim to protect all data, not just the most critical.

While it may minimize disruption to business operations, backing up all data only once per month could result in significant data loss in the event of a disaster or system failure. An effective backup solution should aim for more frequent backups to minimize the risk of data loss.

q_bkp_stor_enterprise-backup_02_secp8

You are the IT manager of a large corporation. Your company has been using a traditional backup method that backs up all data once per week.

However, with the growth of the company and the increase in data, this method is causing performance issues and proving insufficient. You are considering implementing an enterprise backup solution.

Which of the following factors should be your top priority when choosing a new backup solution?

Answers:

The scalability of the backup solution.

***The ability of the backup solution to minimize performance disruptions during backup operations.**

The brand popularity of the backup solution.

The usability of the backup solution interface.

Explanation:

The ability of the backup solution to minimize performance disruptions during backup operations should be the top priority in this scenario. The current backup method is causing performance issues, so it's crucial to choose a solution that can perform backups efficiently without significantly impacting system performance.

While scalability is an important factor to consider when choosing a new backup solution, it is not the top priority in this scenario. The main issue at hand is the performance disruptions caused by the current backup method.

While a popular brand may indicate a reliable and well-tested product, it is not the top priority in this scenario. The specific features and capabilities of the solution and how well they meet the company's needs should be the focus, particularly the need to minimize performance disruptions.

The usability of the backup solution interface does not impact the functionality or effectiveness of the solution and should not be a top priority. The main issue at hand is the performance disruptions caused by the current backup method.

q_bkp_stor_filesystem-snapshots_secp8

A large organization is implementing a data protection strategy and considering the use of snapshots. The company has a complex IT environment with a mix of physical servers, virtual machines, and cloud-based services.

The primary concern is the ability to recover individual files or previous versions of files in case of accidental deletion or data corruption.

Which type of snapshot would be MOST suitable for this organization?

Answers:

Virtual Machine (VM) snapshots

***Filesystem snapshots**

Storage Area Network (SAN) snapshots

Cloud-based snapshots

Explanation:

Filesystem snapshots is the correct answer. Filesystem snapshots, like those provided by ZFS or Btrfs, capture the state of a file system at a given moment. This allows users to recover accidentally deleted files or restore previous versions of files in case of data corruption. This makes them the most suitable option for the organization's primary concern.

While VM snapshots capture the state of a virtual machine, including its memory, storage, and configuration settings, they are not the best option for recovering individual files or previous versions of files. They are more suitable for rolling back the entire VM to a previous state in case of failures or during software testing.

SAN snapshots capture the state of the entire storage volume, allowing for rapid recovery of large datasets and applications. However, they are not as granular as filesystem snapshots and may not be the best option for recovering individual files or previous versions of files.

Cloud-based snapshots are similar to VM snapshots but are taken in a cloud environment. They capture the state of a cloud-based virtual machine or storage volume. While they can be useful in a cloud environment, they may not provide the granularity needed for recovering individual files or previous versions of files.

q_bkp_stor_remote-journaling_secp8

A large multinational corporation is implementing a data protection strategy. The company has a complex IT environment with a mix of physical servers, virtual machines, and cloud-based services.

The primary concern is the ability to track and monitor data modifications and revert to previous states if necessary, especially in the event of local failures, natural disasters, or malicious attacks.

Which type of journaling would be MOST suitable for this organization?

Answers:

File system journaling

Database journaling

***Remote journaling**

Application-level journaling

Explanation:

Remote journaling is the correct answer. Remote journaling creates and maintains a journal of data changes at a separate, remote location, allowing for data recovery and ensuring business continuity in case of local failures, natural disasters, or malicious attacks. This makes it the most suitable option for the organization's primary concern.

File system journaling, such as the Journaled File System (JFS) or the New Technology File System (NTFS), records all changes made to files, allowing for data recovery and consistency checks after unexpected system shutdowns or crashes. While it is useful for tracking and monitoring data modifications, it may not provide the necessary protection in the event of local failures, natural disasters, or malicious attacks.

Database journaling records changes to data within a database. While it can be useful for tracking and monitoring data modifications within a database, it may not provide the necessary protection in the event of local failures, natural disasters, or malicious attacks.

Application-level journaling records changes to data within a specific application. While it can be useful for tracking and monitoring data modifications within a specific application, it may not provide the necessary protection in the event of local failures, natural disasters, or malicious attacks.

q_bkp_stor_replication-secp8

In the context of data protection and recovery, which of the following statements about replication is correct?

Answers:

Replication is a data compression technique that optimizes storage space by identifying and eliminating redundant data.

Replication involves creating and maintaining different versions of data on the same storage system.

Replication can increase the risk of data loss due to hardware failures, human errors, or malicious attacks.

***Replication involves creating and maintaining exact copies of data on different storage systems or locations.**

Explanation:

Replication involves creating and maintaining exact copies of data on different storage systems or locations. Replication is a data protection method that involves creating and maintaining exact copies of data on different storage systems or locations.

Replication is a data compression technique that optimizes storage space by identifying and eliminating redundant data is incorrect. The process described here is data deduplication, not replication. Replication involves creating and maintaining exact copies of data, not eliminating redundant data.

Replication involves creating and maintaining different versions of data on the same storage system is incorrect. Replication involves creating and maintaining exact copies of data on different storage systems or locations, not different versions of data on the same storage system.

Replication can increase the risk of data loss due to hardware failures, human errors, or malicious attacks is incorrect. By having redundant copies of the data, replication can actually safeguard against data loss due to various factors including hardware failures, human errors, or malicious attacks.

q_bkp_stor_san-snapshots_secp8

A large multinational corporation is implementing a data protection strategy. The company has a complex IT environment with a mix of physical servers, virtual machines, and cloud-based services.

The primary concern is the ability to rapidly recover large datasets and applications in the event of a major system failure or data corruption.

Which type of snapshot would be MOST suitable for this organization?

Answers:

Virtual Machine (VM) snapshots

Filesystem snapshots

***Storage Area Network (SAN) snapshots**

Cloud-based snapshots

Explanation:

Storage Area Network (SAN) snapshots is the correct answer. SAN snapshots are taken at the block-level storage layer within a storage area network. They capture the state of the entire storage volume, allowing for rapid recovery of large datasets and applications. This makes them the most suitable option for the organization's primary concern.

VM snapshots capture the state of a virtual machine, including its memory, storage, and configuration settings. While they can be useful for rolling back the entire VM to a previous state in case of failures or during software testing, they may not be the most efficient for rapid recovery of large datasets and applications.

Filesystem snapshots, like those provided by ZFS or Btrfs, capture the state of a file system at a given moment. They are excellent for recovering accidentally deleted files or restoring previous versions of files in case of data corruption. However, they may not be the best option for rapid recovery of large datasets and applications.

Cloud-based snapshots are similar to VM snapshots but are taken in a cloud environment. They capture the state of a cloud-based virtual machine or storage volume. While they can be useful in a cloud environment, they may not provide the rapid recovery of large datasets and applications that SAN snapshots can offer.

q_bkp_stor_vm-replication_secp8

A large software development company is implementing a data protection strategy. The company heavily relies on virtual machines for their development and testing environments.

The primary concern is the ability to quickly recover and resume work in the event of a primary VM failure or corruption.

Which data protection method would be MOST suitable for this organization?

Answers:

File system journaling

Database mirroring

SAN replication

***VM replication**

Explanation:

VM replication is the correct answer. VM replication creates and maintains an up-to-date copy of a virtual machine on a separate host or location, ensuring that a secondary VM can quickly take over the workload in the event of a primary VM failure or corruption. This makes it the most suitable option for the organization's primary concern.

File system journaling records all changes made to files, allowing for data recovery and consistency checks after unexpected system shutdowns or crashes. While it is useful for tracking and monitoring data modifications, it may not provide the quick recovery of a virtual machine environment that the company needs.

Database mirroring involves maintaining primary and secondary mirrored databases. Any changes made to the primary database are automatically replicated to the secondary database. While it ensures data consistency and availability if the primary database encounters any issues, it may not provide the quick recovery of a virtual machine environment that the company needs.

SAN replication duplicates data from one SAN to another in or near real-time, providing redundancy and protection against hardware failures, human errors, or data corruption. While it is beneficial for large datasets and applications, it may not provide the quick recovery of a virtual machine environment that the company needs.

10.0 Protocol, App, and Cloud Security

10.1 Host Virtualization

As you study this section, answer the following questions:

What is virtualization?

What is the difference between a virtual machine and a hypervisor?

What are the advantages of virtualization?

How do you secure a container?

In this section, you will learn to:

Use VMWare Player.

Use Hyper-V.

Create virtual machines.

Use Windows Sandbox.

Create containers.

Secure containers.

Key terms for this section include the following:

Term	Definition
Physical machine	The physical computer with hardware, such as the hard disk drive(s), optical drive, RAM, and motherboard.

Virtual machine	A software implementation of a computer that executes programs like a physical machine.
Virtual hard disk (VHD)	A file that is created within the host operating system and simulates a hard disk for the virtual machine.
Hypervisor	A thin layer of software that resides between the guest operating system and the hardware. It creates and runs virtual machines.
Load balancing	A technique that disperses a workload between two or more computers or resources to achieve optimal resource utilization, throughput, or response time.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
TestOut Security Pro	3.3 Implement virtualization 3.3.1 Create virtual machines 3.3.2 Create virtual switches 5.2 Assessment techniques 5.2.3 Scan for vulnerabilities
CompTIA Security+ SY0-701	2.3 Explain various types of vulnerabilities. Virtualization Virtual machine (VM) escape Resource reuse 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Segmentation 3.1 Compare and contrast security implications of different architecture models. Architecture and infrastructure concepts Network infrastructure Logical segmentation Containerization

	Virtualization 3.2 Given a scenario, apply security principles to secure enterprise infrastructure. Infrastructure considerations Load balancer 3.4 Explain the importance of resilience and recovery in security architecture. High availability Load balancing vs. clustering 4.1 Given a scenario, apply common security techniques to computing resources. Sandboxing 4.3 Explain various activities associated with vulnerability management. Identification methods Vulnerability scan

10.1.1 Host Virtualization Overview (Lesson Video)

Transcript:

Let's spend a few minutes talking about virtualization and virtual network components.

Virtualization is a very important aspect of modern computing. But before we talk about what virtualization is and how it works, we first need to look at the way things have traditionally been done.
In the traditional computing model, we have one operating system, or OS, installed on one hardware device. In this example, we have hardware that consists of our CPU, storage, and memory resources. We need an operating system to run these resources, which could be Windows, Linux, or even macOS. Either way, we have one hardware device being managed by one operating system.
The operating system's role is to provide services for all the applications that run on the hardware device so that they can utilize the memory, storage, and CPU resources efficiently. Here we have App1, App2, App3, and App4. Each app makes its own requests for hardware resources, and then the OS makes that happen. For example, let's suppose you're going to install a new server in your data center. The first thing you'd do is order your server hardware and then you'd pick an operating system. When the server comes, you'd install a server operating system on the hardware. The operating system has full reign over all the resources available on the server hardware, including RAM, processor time, storage devices, network interfaces, and so on. Finally, you'd install applications on top of that operating system that would interact with system resources.
There's a fundamental problem with doing things this wayâ€”it's terribly inefficient. Usually, applications that run on the OS don't fully utilize all the hardware resources all the time. This is often true on a server system, and you find out that a lot of times your computing resources are actually underutilized. For example, you may see your CPU utilization hovering around 10% most of the time. You may also see that your memory utilization is likewise hovering at around 20%. This means 80-90% capacity is still available to perform additional computing tasks.

Servers are expensive, and wasted capacity means wasted dollars. Many times, servers are purchased for a specific purpose such as running web servers for individual apps. This can lead to server sprawl for the many underutilized servers This sucks up lots of power, puts out tons of heat, and costs an immense amount of money.
But there's a way to overcome this and consolidate all those servers into just a few, compact ones. That's where virtualization comes in.
Virtualization enables you to run multiple servers on a single physical server, called a host, through hardware virtualization. When you do this, each independent operating system or virtual machine thinks it has all the hardware resources it needs and runs in its own little isolated container. So instead of the many underutilized servers in traditional networking, with virtualization we can put many servers together on one hardware chassis to fully utilize CPU, RAM, disk, and network subsystems.
Virtualization increases IT agility, allows for greater flexibility, delivers increased performance, and allows you to scale up quickly while enjoying significant cost savings. We'll get to some more important benefits in just a minute.
Nowadays, we can virtualize just about anythingâ€”operating systems, storage, apps, networking, desktops, and more.
Virtualization can also be enabled in the cloud where much of the computing can be done remotely without you having to build and maintain your own data center. It also allows companies to build their own private clouds in their data centers to virtually deliver their homemade apps to their corporate users. Now users can use just about any device to securely access their workloads remotely at any place, anytime, anywhere. The power of virtualization is driving an immense, global, and increasingly mobile workforce.
Virtualization starts with you installing a thin piece of software called a hypervisor, which decouples the virtual machines, or VMs, from the hardware and dynamically allocates resources as the virtual machines request them. Once the hypervisor is installed, you can then create virtual machines with their own operating system and apps.
Notice in this diagram we have our memory, storage, and CPU just like before. But instead of installing an operating system directly onto the host, we install a hypervisor to manage it.
On top of the hypervisor, we create virtual machines by allocating CPU, RAM, and storage. Generally, the more hardware you have on the host, the more virtual machines you can build. These physical servers that host multiple virtual machines have lots of processors and lots of RAM, and they're usually connected to an immense storage area network, or SAN, that's made to be fault tolerant.
Notice that each VM has its own OS and applications in an isolated container. Each VM makes calls to use each piece of hardware, and the hypervisor facilitates those requests. You can install just about any OS inside of a VM and then run apps built for that OS.
Before we go on, there are three kinds of hardware virtualization that you need to be familiar with.
The first type is called full virtualization. In full virtualization, the virtual machine completely emulates a real, physical host. This is good because it allows most operating systems and applications to run within that virtual machine without being modified in any way. The hypervisor fully decouples the operating system from the hardware. There are two types of full virtualization: bare-metal hypervisors (type 1) and hosted hypervisors (type 2). We'll learn more about these in just a minute.
Another form of virtualization is called partial virtualization. In partial virtualization, some virtual machine components are virtualized while others aren't. In other words, the OS uses some virtual components and some physical components while it runs. It's important to note that an operating system can only run in a partial virtualization environment if that operating system has been specifically modified to do so.
The last type of virtualization is called paravirtualization. In paravirtualization, the hardware isn't virtualized at all. In other words, all the guest operating systems running on the hypervisor actually access various hardware resources in the physical devices directly. It doesn't virtualize any of the components. To do this, the guest operating systems run in isolated domains on the same physical hardware, which basically simulates running on separate systems. For an operating system to run on a paravirtualization environment, it must be specially modified to do so.
As you can see, a hypervisor is a fundamental component in any virtualization implementation.
There are several different hypervisors that you can choose from, and they generally fall within two main categories: type 1 and type 2.
As we mentioned, type 1 hypervisors are installed directly onto the hardware and allow you to build VMs. This type is used in data centers to virtualize servers. Examples of type 1 hypervisors include VMware vSphere, Microsoft Hyper-V, Linux KVM, Citrix Hypervisor, Xen, and Oracle VM. These are enterprise-class hypervisors for serious, large-scale virtualization deployments that provide an array of management tools and feature-rich benefits.
Alternatively, we have type 2 hypervisors. Since type 2 hypervisors are hosted and run on top of a regular operating system, they allow you to get a test environment up and running quickly without having to buy huge systems. Examples of type 2 hypervisors are VMware Workstation, Oracle VirtualBox, and Microsoft Hyper-V which is built into Windows 8.1 and 10.

All these hypervisors perform basically the same role that we talked about before. They allow you to install multiple operating systems on a single physical server.
There are many benefits to implementing virtualization in your environment. The biggest drivers are reduced expenses and increased return on investment. Most virtualization platforms provide redundancy and load balancing, which minimize downtime. Also, unlike traditional computing where you have to go out and buy a physical server, you can spin up a virtual server in a fraction of the time. You can also take server snapshots in case you need to make a server rollback. An entire VM is stored in just a few files, and they can easily be copied or moved as you would any file. Finally, virtualization makes IT tasks much easier to manage and less costly.
With benefits like these, you can see why nearly all companies are using some form of virtualization in their business today.
Before we finish, we need to discuss a possible virtualization attack. Virtual machine escape is one of the security threats an attacker can take advantage of from within a virtual machine. Typically after escaping
out of a guest virtual machine, an exploiter can execute code on the hypervisor or the host operating system. The threat actor can gain access to the primary hypervisor and perform administrative tasks like VM deletion, powering off, resource allocation, and much more. The best way to prevent virtual machine escape is to regularly patch your hypervisor. Keeping your hypervisor current goes a long way in preventing exploits.
That's it for this lesson. In this lesson, you were introduced to virtualization's common roles and functions. We discussed how virtualization works, different virtualization types, the hypervisors you can use, and virtualization benefits. We ended this lesson by discussing virtual machine escape and how you can prevent these attacks.

10.1.2 Load Balancing with Virtualization (Lesson Video)

Transcript:

Virtualization provides many advantages above simply installing an operating system on a physical computer system. One powerful benefit and feature is load balancing. Let's look at how load balancing works with virtual machines, or VMs.

VMs run on top of a hypervisor which is installed on physical hardware called a host. The hypervisor decouples the hardware from the VMs, or workloads.
If we have two or more hosts, we can create a cluster. Clusters usually require shared storage like a SAN where the actual VM files reside and can be managed as a single entity. While workloads are running in memory, they still use real resources â€“ real physical memory, real storage, and really CPUs, and the cluster manages that linking in an effective, optimized way.
For example, if the cluster contains eight host servers with four quad-core CPUs each running at 4GHz and 64GB of memory, the cluster has an aggregate of 512GHz of CPU power and 512GB of memory available for running VMs. Having a cluster provides a number of benefits, including high-availability, load balancing, and high performance. Let's look at load balancing.
A cluster can sense which hosts are busy and offload workloads to a less busy host. It does this by monitoring hardware resources like CPU and RAM on each host of the cluster to determine the best placement for all the workloads. The load is evaluated frequently, and workloads can be adjusted dynamically to reallocate when needed.
This means a VM or workload could move from one host to another in a cluster for better performance. And what's cool is this is seamless to the userâ€”they don't even notice all this is happening on the back end.
This dynamic allocation load balances the workloads across hosts in the cluster based on host resource load and availability. When a host gets busy, workloads will be moved to another host on the fly. That's powerful, and ensures each workload gets the very best performance possible.
Not all your hosts will have the same hardware with CPU and RAM resources. Some hosts will be beefier than others. Some will have newer processors. This means that some hosts can handle more dense workloads than others. The cluster can sense how busy each host is and adjust, or load-balance, the workloads for best performance and to best utilize CPU and memory resources across the cluster.
Within a cluster, we can usually create what's called a resource pool. This helps us allocate CPU and RAM resources to prioritize our workloads.
Some workloads may have applications that are mission critical and require specific resources to always be up and running properly. Other workloads may not be as important, but still need adequate resources. Resource pooling allows us to establish performance boundaries within the cluster, so resources aren't consumed by low-level workloads and keep important workloads from starving for resources.

So, a resource pool works within the cluster, and may span two or more host servers. Let's create Resource Pool A with a certain amount of RAM and CPU shares. After creating the pool, we can then assign or create VMs to run in the pool. We can create as many pools as we like in our cluster, but we must always assign them CPU and memory performance. Here's another pool â€“ Resource Pool B. We'll assign it a certain amount of CPU and RAM, and then we'll assign VMs to it.
In essence, the workloads end up being load balanced, and they are throttled to only use the established resource limits. This allows admins to aggregate computing capacity and delegate it for use rather than focus on individual host servers. And, as a bonus, if one host server goes down, the other servers in the cluster will pick up the workloads and continue on.
You can create a resource pool on an individual host. However, if you have a cluster, resource pooling will provide more resources across the hosts in a cluster, allowing you to better manage the use of those hardware resources, including reservations for CPU and RAM usage. The VMs will be load balanced within the resource pool within the cluster. This provides for rapid elasticity which leverages virtualized server and storage technologies to rapidly meet the rise and fall of user load and service demand.
Hypervisor software that supports resource pooling is downloadable from the internet. You can download an evaluation copy to set up and experiment with load balancing on virtual machines.
That's it for this lesson. Within the context of virtualization, that's how load balancing works. Clustering allows VMs to be load balanced and highly available while providing the best performance possible for each workload. And resource pooling allows us to aggregate hardware resources across multiple hosts ensuring that the workloads get the resources they need without sacrificing the performance of others.

10.1.3 Virtualization Facts

This lesson covers the following topics:

- Virtualization components
- Virtualization advantages and disadvantages
- Virtualization security
- Load balancing

Virtualization Components

Virtualization refers to installing and running multiple operating systems concurrently on a single physical machine. Virtualization typically includes the following components:

Component	Description
Physical machine	A *physical machine* , also known as the host operating system, has the hardware, such as the hard disk drive(s), optical drive, RAM, and motherboard.
Virtual machine	A *virtual machine* , also known as the guest operating system, is a software implementation of a computer. The virtual machine executes programs in the same way a physical machine executes programs. The virtual machine appears to be a self-contained and autonomous system.
Virtual hard disk (VHD)	A *virtual hard disk* (VHD) is a file created within the host operating system and simulates a hard disk for the virtual machine.

Hypervisor	A *hypervisor* is a thin layer of software that resides between the guest operating system and the hardware. A hypervisor allows virtual machines to interact with the hardware without going through the host operating system. There are two types of hypervisors. A Type I hypervisor is often called a native hypervisor or bare-metal hypervisor. A hypervisor in a dedicated appliance is called an embedded hypervisor. A Type I hypervisor is like a thin operating system that directly interfaces with the computer hardware. Examples of Type I hypervisors are: VMware ESX and ESXi Microsoft Hyper-V Linux KVM Citrix Hypervisor Xen Oracle VM A Type 2 hypervisor is known as a hosted hypervisor. It runs as an application on a conventional operating system. While it may be used in a production environment, a type 2 hypervisor is most often used as a development sandbox. Examples of Type II hypervisors are: VMware Workstation and VMware Player Oracle Virtual Box Microsoft Hyper-V built into Windows 8.1 and 10 Parallels Desktop for Mac
Containerization	Containerization is an alternative to using a hypervisor, which enforces resource separation at the operating system level. The OS defines isolated "cells" for each user instance to run in. Each cell or container is allocated CPU and memory resources, but the processes all run through the native OS kernel. These containers may run slightly different OS distributions but cannot run different types of guest OSes (you could not run Windows or Ubuntu in a Red Hat Linux container, for instance). Alternatively, the containers might run separate application processes, in which case the variables and libraries required by the application process are added to the container. One of the best-known container virtualization products is Docker (docker.com). Containerization underpins many cloud services. In particular, it supports microservices and serverless architecture. Containerization is also being widely used to implement corporate workspaces on mobile devices.

Virtualization Advantages and Disadvantages

The advantages of virtualization are described in the following table.

Advantage	Description
Flexibility	Virtual machines can be given network access. Other network devices will consider them to be physical machines. Be aware that virtual machines: Should have the latest service packs and patches, just like physical machines. Should be hardened, just like physical machines. Can be connected to the production network by creating a bridged (external) virtual switch. Because they are self-contained, virtual machines can be easily moved between hypervisor hosts as needed.
Security	To better protect other systems, virtual machines can be used to create honeypots and honeynets to attract attackers so you can analyze attacks on the system.
Testing	Virtual machines can be configured in a lab environment that mirrors a production network. This lab environment can be used to: Test applications before installing them on production systems. Test updates and patches before rolling them out into the production environment. Test security controls to verify that they are working as designed.
Server consolidation	Server consolidation allows you to move multiple physical servers onto a few physical servers with many virtual machines. *Physical-to-virtual migration* (P2V migration) is the process of moving an older operating system off aging hardware and into a virtual machine. Consolidating servers: Requires fewer physical computers. Reduces power consumption. Increases physical server utilization of resources. Increases administrative efficiency. Reduces the number of incompatibility issues.
Isolation	A virtual machine can be isolated from the physical network to allow testing to be performed without impacting the production environment. This is called sandboxing.

	Sandboxed virtual machines offer an environment in which malware can be executed with minimal risk to equipment and software. Sandboxing virtual machines protects them from many kinds of security threats. To allow isolated virtual machines to communicate with each other, create a new virtual switch configured for host-only (internal) networking. Connect the virtual network interfaces in the virtual machines to the virtual switch.
Application virtualization	Applications can be virtualized. A virtual application appears to be local but is really running on a different system. Virtualized browsers can protect the underlying physical operating system from malware installation. Any malware installed from the virtual browser affects only the browser, not the rest of the system. Malware can also use virtualization techniques that make it difficult to detect.

Disadvantages of virtualization include:

Disadvantage	Description
Attacks	An attack on the host machine could compromise all guest machines operating on that host.
Bottleneck	A bottleneck or failure of any hardware component that is shared between multiple guests, such as a failure in a disk subsystem, could affect multiple virtual machines.
Complexity	While administration is centralized, virtualization is a newer technology and requires new skills. Managing virtual servers can add complexity.
Server sprawl	Your configuration is susceptible to server sprawl, a condition that delays patch and security update management due to the number of virtual machines that must be managed.
VM escape	A VM escape is when an attacker with access to a virtual machine breaks out of this isolated environment and gains access to the host system or other VMs running on the same host. Such a vulnerability could allow an attacker to gain control of all virtual machines running on a single physical server, leading to a potentially devastating security breach. A famous example is the "Cloudburst" vulnerability in VMware's virtual machine display function. The Cloudburst vulnerability, officially designated as CVE-2009-1244, was a critical security flaw discovered in 2009 in VMware's ESX Server, a popular enterprise-level virtualization platform. A vulnerability in the virtual machine display function allowed a guest operating system to execute code on the host operating system.

Resource reuse	Virtual machines are frequently created, used, and then deleted in a virtualized environment. If the resources, such as disk space or memory, are not properly sanitized between each use, sensitive data could be leaked between virtual machines. For instance, a new virtual machine may be allocated disk space previously used by another VM. If this disk space is not properly wiped, the new VM could recover sensitive data from the previous VM. Thorough data sanitization practices, ensuring data encryption throughout the lifecycle, and implementing robust encryption key management practices mitigate the risk of resource reuse in cloud infrastructure. Training on cloud provider security features and best practices and segregating resources based on security levels also mitigate risks.
Hypervisor	Attackers exploit hypervisors to gain unauthorized access and compromise the virtual machines (VMs) running on them. Hypervisors typically provide specialized management interfaces so administrators can control and monitor their virtualized environments. These interfaces can become potential attack vectors if insecure. For example, weak authentication, lack of encryption, or vulnerabilities in communication protocols can lead to unauthorized access to the virtualized environment. Like any software, hypervisors have vulnerabilities that must be regularly patched.

Virtualization Security

Security considerations for a virtual machine should be the same as for physical machines. For the host and all guest machines, be sure to:

- Reduce the number of services running.
- Apply patches and updates regularly.
- Install antivirus and other security software.
- Implement backups, operating system snapshots, and other solutions for data protection.

In addition, you should protect against virtual machine escape, an exploit in which malware allows the operating system within a virtual machine to break out and interact directly with the hypervisor. To minimize this vulnerability:

- Apply patches and updates regularly.
- Install only the resources-sharing features that are necessary.
- Install only the software applications that are necessary.

Load Balancing

Load Balancing is a technique that disperses a workload between two or more computers or resources to achieve optimal resource utilization, throughput, or response time. The primary goal of load balancing is to improve performance and create high availability by configuring multiple devices to respond as one. Load balancing can also provide fault tolerance.

If the load balancing mechanism is able to detect when a specific node or member is unavailable, new requests will automatically be distributed to other available members. Load balancing methods with virtualization include the following:

Resource pooling creates shared logical pools of CPU and memory resources from many physical machines within the hypervisor to guarantee a level of resources for specific virtual machines.

Workload balancing distributes a workload (the total requests made by users and applications of a system) across multiple computers or a computer cluster to achieve optimal resource utilization, maximize throughput, minimize response time, and avoid overload.

10.1.4 Use VMWare Player (Demo Video)

Transcript:

In this demo, we'll be working with VMware Workstation Player, which is software that allows the installation of virtual machines, or VMs. This software is similar to commercially available hypervisors, but on a much smaller scale. VMware Workstation Player is designed to be used on your desktop for either personal or business use. Today, we'll just be looking at the standard edition that you'd use in a home lab or for testing.

Let's go ahead and walk through a virtual machine installation. We have several options off to the right. Create a New Virtual Machine will start the wizard to add a virtual machine to our player. Open a Virtual Machine will add a virtual machine that was already created elsewhere or with an open virtual machine format, such as an .ovf file. There's an upgrade option here to get Workstation Pro, which gives you more features than would typically come with VMware's ESXI hypervisor, but this requires a paid license. You can always get help if you're feeling stumped and want to check out more information. We're going to go ahead and create a new VM.
Typically, operating systems have an ISO that's downloadable from their website. We have some ISOs already downloaded, so we're just going to click Browse. We'll select our Win11 ISO file and click Open. From here, we just need to click Next. Quickly make sure that the Guest operating system is set appropriately. In our case, it'll be Microsoft Windows and the version will be Windows 10 and later x64. Click Next.
Here it gives you an option to name the VM and its location. Just to make things easier, I'll change the name to Windows11 and leave the location alone. Click Next. Okay, this is one feature that's sort of unique. We have the ability to store the VM's disk as a single file or break it up into smaller files. It says that it might be easier to move the smaller files, but in our case, we just want to keep it as one. Also, be mindful of how much disk space you have on your workstation and how much you plan to use in the future on this virtual machine. There are usually bare bone requirements for all operating systems. Of course, you may want to add more if there's another purpose to this VM. Let's click Next.
Notice that it didn't give me the option to tweak more VM settings. If you do want to finetune things more, like adding more CPU or memory, tweaking the network, etc., you can customize the hardware here or do it after the fact by right-clicking on the VM and going to Settings. Click Finish.
Now we have a newly created VM ready to boot. Let's go over the Settings that we have. Some of this is self-explanatory, like adding or removing Memory, Processors, or even Hard Disk space. The Network Adapter is where it gets interesting. There are many types of network connections to choose from. Bridged connects directly to the physical network card, NAT shares the host's IP address, Host-only puts the VM in an isolated private network, Custom drops it in a virtual network, and LAN segment does something similar. NAT may be your best choice if your VM just needs to reach the internet and isn't going to be used in conjunction with other VMs. If you want to network several VMs together, you can use a custom virtual network or LAN segment.
For example, if I wanted to create a firewall VM that all my VMs go through first before they hit the internet, I might have two network adapters in here. One of them would be the NAT, the internet source, while the secondary one would be the LAN or a custom virtual network of my choosing. Then all the other VMs would use the virtual network and not the NAT connection.
In our demo, we're just going to leave it on NAT and click OK.
Now it's time to power this on. Clicking Play virtual machine makes it come alive. And just like that, we're booting the Windows 11 ISO to start the installation process. Keep in mind that you have to leave the player open to run these VMs. These aren't meant to run what they call "headless" and in the background with VMware Workstation player. Once this application is closed out, your VMs close out as well.
That's it for this demo on virtual machines. In this demo, we discovered VMware Workstation Player. We also showed you how to install a VM and tweak its settings.

10.1.5 Use Hyper-V (Demo Video)

Transcript:

Hyper-V is virtualization software created by Microsoft in 2016 that lets you run other software like it's on a physical computer. We call this a virtual machine, or VM. With a VM, you're able to have many different operating systems installed on one physical system.

Today, we're going to look at how to enable Hyper-V on a desktop. In the Start menu, look for the features keyword, and it'll direct you to the section to add or remove features. Alternately, you could go to Control Panel > Programs > Turn Windows features on or off. This is the area that controls extra features that come with Windows, ones that might not be enabled by default. Also, before we go on, it's important to know that you won't be able to enable Hyper-V unless you have the Pro version of Windows.
Now, after clicking OK, it'll process your request when it's done and require a reboot. We're going to do that now, so I'll pause the recording.
Okay, our reboot is complete, and we're ready to launch Hyper-V. To do so, we go to Search and type hyper-v. This brings us the Hyper-V Manager. Eventually, when you have VMs, they'll be populated over here. This manager allows you to control all the local Hyper-V instances along with remote ones.
We're only worried about our local one, so we're going to right-click here and select Virtual Switch Manager. Here's where we manager the virtual switches that'll be used for the VMs. It's a good idea to configure these first so you know how your network is going to be set up. And there are three options to select from here. External binds to the physical network adapter. This would give a VM access to the internet and the physical network that your PC is connected to. The Internal switch is similar except that it doesn't allow direct access to the internet. This would be a good option if you had several VMs that needed to talk to each other without access to a physical network. Private is completely isolated. Some use cases for this would be restoring a VM for testing purposes that you don't want to communicate with anything.
We're going to do something simple and select Internal. Let's give it a name of Internal as well. We'd have the ability to select the Ethernet controller or Wi-Fi adapter if we'd used the external network switch. If there were other network adapters available, they'd be listed here, too. And if you were running VLANs, you could select this and tag the network with a certain VLAN ID. Click Apply when you're done creating your switch.
One thing to notice here is the Default Switch. This switch does use the Internal network setting, but the part at the bottom makes it different from the internal network we just created. This has NAT, or Network Address Translation, set up to give your VMs access to the computer's physical network. If the NAT weren't on, you wouldn't have access to the physical network. This default is set this way to make it easy for users to start right out of the box. Anyway, click OK.
Now we're ready to create our first VM. Right-click on the host and select Quick Create. In here, we select the ISO of our choosing from some default options, or we could just upload our own from a local source. I happen to have one in my Downloads folder, so I'm going to go search and grab my Windows ISO. After selecting Open, we're able to go through and create our virtual machine.
Before booting, we need to check out some settings first. They make it handy by giving us the ability to launch the settings from here. In here are all the settings related to the VM we just created. We're able to tweak quite a few aspects, like Memory, Processor, Hard Drive, and more. Let's adjust some of these. Our memory is a little low, for default specs of Windows, we can up it to 4096. It does have the ability to dynamically change as well. Under Processor, it's set to 8. Let's adjust it to 2. Under Network Adapter, we need to configure this VM to use the network we created. For our case, we need to use the isolated network. However, it's possible to create a firewall VM with two network adapters, one being connected to the physical network and the second one to the internal network we just created. Essentially, that'd give you total control of what comes in and out of your Hyper-V host VM. Lastly, let's give this a name. This could be a convention or something of your choosing. Apply and click OK.
Connect allows us to start the process to power on the VM. This is essentially the VM's console, just like you'd see from a computer booting up from a monitor. Click Start. And just like that, we've started the process to install Windows.
And that's it for this demo. In this demo, we went over how to install Hyper-V. We also configured a new virtual switch and created a VM.

10.1.6 Create Virtual Machines (Simulation)

Scenario

You have installed Hyper-V on ITAdmin. You're experimenting with creating virtual machines.

In this lab, your task is to create two virtual machines named VM1 and VM2. Use the following settings as specified for each machine:

VM1:

- Virtual machine name: **VM1**
- Virtual machine location: **D:\HYPERV**
- Generation: **Generation 1**
- Startup memory: **1024 MB** (do not use dynamic memory)
- Networking connection: **External**
- Virtual hard disk name: **VM1.vhdx**
- Virtual hard disk location: **D:\HYPERV\Virtual Hard Disks**
- Virtual hard disk size: **50 GB**
- Operating system will be installed later

VM2:

- Virtual machine name: **VM2**
- Virtual machine location: **D:\HYPERV**
- Generation: **Generation 1**
- Startup memory: **2048 MB** (use dynamic memory)
- Networking connection: **Internal**
- Virtual hard disk name: **VM2.vhdx**
- Virtual hard disk location: **D:\HYPERV\Virtual Hard Disks**
- Virtual hard disk size: **250 GB**
- Operating system will be installed later
- Minimum RAM: **512 MB**
- Maximum RAM: **4096 MB**

Explanation

While completing this lab, use the following virtual machine (VM) specifications:

VM1:

Virtual machine name: **VM1**

Virtual machine location: **D:\HYPERV**

Generation: **Generation 1**

Startup memory: **1024 MB** (do not use dynamic memory)

Networking connection: **External**

Virtual hard disk name: **VM1.vhdx**

Virtual hard disk location: **D:\HYPERV\Virtual Hard Disks**

Virtual hard disk size: **50 GB**

Operating system will be installed later

VM2:

Virtual machine name: **VM2**

Virtual machine location: **D:\HYPERV**

Generation: **Generation 1**

Startup memory: **2048 MB** (use dynamic memory)

Networking connection: **Internal**

Virtual hard disk name: **VM2.vhdx**

Virtual hard disk location: **D:\HYPERV\Virtual Hard Disks**

Virtual hard disk size: **250 GB**

Operating system will be installed later

Minimum RAM: **512 MB**

Maximum RAM: **4096 MB**

Complete this lab as follows:

Access the Hyper-V Manager.

Select **Start** .

Expand **Windows Administrative Tools** and then select **Hyper-V Manager** .

Create virtual machines on ITAdmin.

Use all default settings unless directed otherwise.

Right-click **ITADMIN** and then select **New** > **Virtual Machine** .

From the Before You Begin dialog, select **Next** .

In the Name field, enter ***VM_ name*** and then select **Next** .

Make sure **Generation 1** is selected, and then select **Next** .

In the Startup memory field, enter ***size*** .

Set the **Use Dynamic Memory for this virtual machine** appropriately, and select **Next** .

Use the Connection drop-down menu to select ***connection_type*** and then select **Next** .

In the Size field, enter ***disk_size*** and then select **Next** .

Make sure **Install an operating system later** is selected and then select **Next** .

Review your configuration and then select **Finish** to create the virtual machine.

Repeat step 2 to create the second virtual machine.

Adjust virtual machine memory for VM2.

From the Hyper-V Manager, under Virtual Machines, right-click **VM2** and select **Settings** .

From the left pane, select **Memory** .

In the Minimum RAM field, enter **512** .

In the Maximum RAM field, enter **4096** .

Select **OK** .

10.1.7 Use Windows Sandbox (Demo Video)

Transcript:

As an IT administrator, you'll often be asked to install new software on the Windows systems running in your office. Before installing any new software you're unfamiliar with, it's always important to first test the software to see how it will function. This can be done using Hyper -V virtual machines, but that requires the extra steps of installing the operating system and getting things ready. It also requires additional space on your disk to store the virtual machine.

To make testing software easier, Microsoft offers a lightweight desktop environment known as Windows Sandbox, where you can safely run applications in an isolated environment. The beauty of using Windows Sandbox is that once Windows

Sandbox is closed, all the software and files that were added when the application was installed are deleted. This means you get a brand-new instance of the sandbox every time you open the application.
To use Windows Sandbox, you must add this feature to your system.
To do this, first, open the Settings app. Next, select Apps. Now, from the right, select Optional features. From this page, scroll down and select More Windows features. It's from here that you can add or remove features.
From this dialog, scroll down and select Windows Sandbox and click OK. If the Windows Sandbox option is unavailable, your computer doesn't meet the requirements to run Windows Sandbox. As you can see, after the feature has been installed, you'll be required to restart your system, so I'll click Restart now.
With the computer restarted, you can now use Windows Sandbox to install and test a new application. To do this, let's click on Start and select All apps. I'll then go down and run Windows Sandbox. Notice that the Windows Sandbox window is opened, which looks like another Windows machine.
For this demo, I'm going to install the Audacity program. To save time, I've already downloaded the install file for this program and saved it to my Downloads folder on the host computer. Let's copy and paste this to the desktop of our sandbox. Now, let's maximize the sandbox window for easier viewing.
From this point, you simply install the application as you would in a normal Windows environment. So, let's double-click on the install file and follow the remaining prompts to install Audacity. Once the application has been installed, I can test it to verify it didn't have any ill effects on my system and functioned the way I wanted it to.
When you're done testing your application, all you need to do is close Windows Sandbox. Notice that when you do, you're warned that all your changes will be permanently lost. That's fine, so I'll click OK.
Just to verify that everything was deleted, let's open Windows Sandbox again. This time let's use the Search feature to find and open this tool. Notice that the installation file for Audacity is gone, as well as the shortcut to run the program. And when I access All apps, you can see that Audacity isn't available.
That's it for this demonstration. In this demo, we showed you how to install the Windows Sandbox feature. We then showed you how to use Windows Sandbox to install and test applications.

10.1.8 Create Containers (Demo Video)

Transcript:

Were going to give you an overview on how you can create a container. If your not familiar with a container it is similar to a virtual machine however instead of having the overhead of several operating systems for each software you can just run the software from an image a container engine. Today were going to be using a container engine called docker desktop. This is widely used by developers and for the use of learning containers. We already have this installed so lets go ahead and launch it from the start menu.

Typically to start you will have an image in which you would build a container from. If you were a developer you would of course make your own image with the required software in it to run as a container. If you don't have an image and want to start from a basic one you could go up to the search bar and type apache for example. If you don't know what apache is it is a web server. In order to grab an image we need to pull it. Once we click pull it will download the image and put it in our Image repository. If there are any updates to an image you would have to pull again to obtain your update.
Now that we have our image downloaded we can now create a container from the image. By clicking on the run option from the image it starts the process to create the container. It does give some optional settings like changing the ports or the volumes it may use. Were going to assign it a random port so we will just type zero here and click run. To check our running containers we can venture to the other menu called containers. All running containers will be displayed here. You can see we have one running which is our web server. Locally on the container it is using port 80 however its redirecting it to port "" . clicking the port reference will allow us to launch the web page from our browser. Since we actually see a page and its presenting a text means our container is working properly. Not all containers will have a web page like this one, for example if its a database container you will have connecting ports available but no website to test.
Along with the gui you can still utilize command line tools to manage your docker images and containers. Were just going to check the status of our running container by launching our terminal. Once launched we can type the command "docker ps" this will show us all our running containers. Once we have the container we want we can run a command like "docker stop c764f54a8429" to stop a container. You can see from the docker ps the container isnt running anymore. Each container has a unique container ID. This ID has to be used when using these commands to manage a container. Just like we can run "docker start c764f54a8429" to start the container we just stopped. Running a docker ps will show the container is running again.

That is it for this demo. In this demo we showed you how to create a container on docker desktop. We also showed you docker containers can be managed by command line tools.

10.1.9 Secure Containers (Demo Video)

Transcript:

In this demo, we will explore various ways to enhance the security of your containers.

In the world of IT, we are always vigilant, searching for vulnerabilities that hackers might exploit to compromise our systems. This concern isn't limited to just bare-bone servers and virtual machines; it extends to containers as well. Containers, such as this Apache web server, are streamlined versions of Linux with only the essential software required to run Apache. When we inspect the image we have in our image store, we can see that it has 32 vulnerabilities with the current container we are running. The top one is marked as "high" priority. The medium and low vulnerabilities are less concerning. If you wish to learn more about the Common Vulnerabilities and Exposures (CVE) or Debian Security Advisories (DSA) associated with these vulnerabilities, you can access that information through the provided links. You can also observe the breakdown of the Docker image, which consists of multiple layers, each possibly having vulnerabilities.
Now, if you examine the image closely, you'll find it comprises two parts: a base Debian image and an httpd image. Keep this in mind as we explore potential fixes. Notably, the httpd image doesn't have any updates, whereas the base image does. While you can review the recommended base image, remember that the image we downloaded can only be modified by its owner. As we are not the owner, our best course of action is to either pull down a new image or create our own image to address this issue. Let's start by attempting to pull a new image using the option to "pull a new image." However, as you can see, when we exit and re-enter, there are no changes to the vulnerabilities.
One potential solution is to update the running container to patch vulnerabilities. Keep in mind that if this container were deleted and rebuilt from the image, the changes we make inside the container will be lost. The primary aim is to address vulnerabilities while waiting for an updated image. As we have one running container from our Apache container image, let's proceed. Within the exec window, we can use the command line to access the Linux system inside the container. Running "apt update" will download the latest packages from the repository, similar to a regular Linux server. If you are wondering if the high vulnerability fix is included in these packages, you can check by running "apt list --upgradable" to see if it's on the list. It appears it isn't, but there are other fixes we can apply. Running "apt upgrade" will begin the process to upgrade the packages. Let's proceed and confirm. We'll fast forward a bit.
Now our container is patched. It's a good practice to restart the container after applying patches. As mentioned earlier, this is a workaround until the image is updated with the latest patches.
Each running container will receive an IP address. The network configuration may vary depending on where your containers are hosted. In our specific setting under "resources" and "network," you can observe a Docker subnet of 192.168.65.0/24, which is the default setting. However, you may need to adjust this based on your requirements. The core concept behind network segmentation is to grant designated servers and users access only to specific network resources. For example, if our accounting department uses an Apache website container, while the marketing department does not, we can block the marketing department from using it. Adding more hurdles within the network may make it more complex, but it also reduces the attack surface for potential hackers.
Simple adjustments, such as using alternate ports, can enhance security. If you examine our currently running container, you'll see it might be using port 80, but it's being redirected to port number 32771. The idea behind using alternate ports is that hackers often scan for specific port numbers. Typically, these are the default port numbers like 80 or 443 for HTTP and HTTPS. While you can't change port 443 for a regular website on the internet, you can do so on your private network. For instance, if we were to set up another Apache website, we could run it on a different port. In the optional settings, you can either allow it to select a random port or choose one yourself. Let's select "8080" as an alternate port to port 80. Now, if we try to launch the web page, you can see it works on port 8080. It's worth noting that this container now requires the patches applied to the previous container.
There are various methods to enhance the security of your containers, and these are just a few strategies to consider when securing your network. That concludes this demonstration. In this demo, we discussed how to identify vulnerabilities within a container image. Additionally, we explored the concepts of network segmentation and the use of alternate ports for added security.

10.1.10 Practice Questions (Section Quiz)

q_virt_browser_01_secp8

You have a development machine that contains sensitive information relative to your business. You are concerned that spyware and malware might be installed while users browse websites, which could compromise your system or pose a confidentiality risk.

Which of the following actions would BEST protect your system?

Answers:

***Run the browser within a virtual environment.**

Configure the browser to block all cookies and pop-ups.

Run the browser in protected mode.

Change the security level for the internet zone to High.

Explanation:

To best protect your system, run the browser in a virtual environment. Virtualization creates an environment that is logically separated from the main system. Any problems that occur within the virtual environment are contained within that environment and do not affect the rest of the system.

While configuring the browser to block all cookies and pop-ups can somewhat help, in this scenario, you need a more robust solution, such as running the browser within a virtual environment.

While running the browser in a protected mode limits how much access malware, spyware, or other potentially harmful code has to your system, it does not work as well as running the browser within a virtual environment.

Once again, setting the internet security zone to High can provide some measure of security; in this scenario, the best solution is to run the browser within a virtual environment.

q_virt_browser_02_secp8

Which of the following is an advantage of a virtual browser?

Answers:

***Protects the host operating system from malicious downloads**

Prevents adware and spyware that monitor your internet activity

Prevents phishing and drive-by downloads

Filters internet content based on ratings

Explanation:

A virtual browser operates within a security sandbox that keeps activities within the browser from affecting the rest of the system. For example, malware downloaded by the virtual browser is limited to the security sandbox and cannot harm the operating system.

The virtual browser does not prevent adware, spyware, or phishing. These threats are still possible within the virtual browser. However, if malware is installed within the virtual session, the malware cannot harm the rest of the system, and the virtual browser can be easily restored to remove the malicious software.

q_virt_containerization_01_secp8

To address the escalating operational costs and complexities stemming from multiple standalone applications, an organization plans to restructure its software deployment process.

They want to minimize overhead, increase flexibility in development environments, and enhance the efficient use of system resources.

What approach would be the MOST effective?

Answers:

***Containerization**

Virtualization

Microservices

Hybrid cloud infrastructure

Explanation:

Containerization encapsulates applications and their dependencies, which provides the flexibility to run them across different environments. It also allows for better resource utilization, as each container only holds the application and its related binaries/libraries, effectively reducing overhead.

Virtualization can provide some cost savings and flexibility but generally consumes more resources than containerization, as each virtual machine requires its own operating system.

Microservices bring flexibility and go in a container, but they do not directly lead to better resource utilization, as they do not impact how applications use system resources.

Hybrid cloud infrastructure provides scalability and potential cost savings but does not enhance the efficient use of current system resources.

q_virt_containerization_02_secp8

A tech company wants to increase its security measures by isolating its various development, testing, and production environments.

The company wants to ensure that these environments are reproducible and managing these dependencies consistently.

Which approach would be MOST beneficial in meeting these requirements?

Answers:

***Containerization**

Virtualization

Software-defined networking (SDN)

Industrial control systems (ICS)/supervisory control and data acquisition (SCADA)

Explanation:

Containerization allows the isolation of environments while also providing reproducibility and efficient dependency management. Each container includes the application and all its dependencies, ensuring that it will run consistently across different environments.

Virtualization does not inherently provide the level of isolation between applications or manage dependencies as effectively as containerization.

Software-defined networking (SDN) can be useful for managing network traffic and isolating parts of the network, but it does not address the reproducibility and dependency management concerns mentioned in the scenario.

Industrial control systems (ICS)/supervisory control and data acquisition (SCADA) systems control industrial processes and collect data from them, respectively. They are not particularly relevant for creating isolated and reproducible software development environments.

q_virt_escape_01_secp8

Which of the following is an exploit in which malware allows the virtual OS to interact directly with the hypervisor?

Answers:

***Escape**

Jump

Bottleneck

Load balancing

Explanation:

Virtual machine escape is an exploit in which malware allows the operating system within a virtual machine to break out and interact directly with the hypervisor.

Jump is not a type of VM exploit.

Load balancing is a technique that disperses a workload between two or more computers or resources to achieve optimal resource utilization, throughput, or response time.

A bottleneck is an area (software, hardware component, etc.) where all traffic slows down.

q_virt_escape_02_secp8

Evaluating the security risks tied to a company's virtualized environment, a network administrator relies on multiple virtual machines (VMs) for service distribution throughout the infrastructure.

The administrator's primary concern is a specific risk: an application compromised within a VM might interact with the host system and potentially propagate harmful activities.

What strategy should the network administrator employ to mitigate this risk and guarantee the individual VMs' integrity and isolation in the company's virtualized environment?

Answers:

***Mitigating virtual machine escape**

Implementing network segmentation

Establishing firewall rules

Enforcing user access controls

Explanation:

Virtual machine escape is a critical security issue in a virtualized environment, where an application, process, or user within a virtual machine can bypass the virtual machine's barriers and interact directly with the host system.

Network segmentation enhances security but does not directly address a virtual machine escape scenario.

Firewalls mainly govern inbound and outbound network traffic. They do not prevent a virtual machine escape situation.

Enforcing user access controls is primarily about controlling who has access to specific resources within a system or network. It does not directly secure a VM from the possibility of breaking its isolation and interacting with the host system.

q_virt_hyper_secp8

Which of the following devices is computer software, firmware, or hardware that creates and runs virtual machines?

Answers:

Virtual switch

Virtual router

Virtual firewall

***Hypervisor**

Explanation:

A hypervisor is computer software, firmware, or hardware that creates and runs virtual machines. A computer on which a hypervisor runs one or more virtual machines is called a host machine. Each virtual machine is called a guest machine. The hypervisor provides the guest operating systems with a virtual operating platform and manages the execution of the guest operating systems.

Virtual switches, routers, and firewalls are part of the virtual environment that enable transmission of network data but have nothing to do with creating and running the virtual machines.

q_virt_load_balance_secp8

Which of the following is a technique that disperses a workload between two or more computers or resources to achieve optimal resource utilization, throughput, or response time?

Answers:

Virtualization

Hypervisor

***Load balancing**

Bottleneck

Explanation:

Load balancing is a technique that disperses a workload between two or more computers or resources to achieve optimal resource utilization, throughput, or response time. The primary goal of load balancing is to improve performance and create high availability by configuring multiple devices to respond as one.

A hypervisor is a thin layer of software that resides between the guest operating system and the hardware.

Virtualization refers to installing and running multiple operating systems concurrently on a single physical machine.

A bottleneck is an area (software, hardware component, etc.) where all traffic slows down.

q_virt_sandbox_secp8

What is isolating a virtual machine from the physical network to allow testing to be performed without impacting the production environment called?

Answers:

Resource pooling

***Sandboxing**

Testing

Workload balancing

Explanation:

Isolating a virtual machine from the physical network to allow testing to be performed without impacting the production environment is known as sandboxing.

Resource pooling creates shared logical pools of CPU and memory resources from many physical machines within the hypervisor. This guarantees a level of resources for specific virtual machines.

Virtual machines can be configured in a lab environment that mirrors a production network to provide a testing environment.

Workload balancing distributes a workload (the total requests made by users and applications of a system) across multiple computers or a computer cluster to achieve optimal resource utilization, maximum throughput, minimal response time, and less overload.

q_virt_server_secp8

Which of the following is a disadvantage of server virtualization?

Answers:

***A compromised host system might affect multiple servers.**

It increases hardware costs.

A compromised guest system might affect multiple servers.

Systems are isolated from each other and cannot interact with other systems.

Explanation:

Virtualization allows a single physical machine (known as the host operating system) to run multiple virtual machines (known as guest operating systems). The virtual machines appear to be self-contained and autonomous systems. Disadvantages of virtualization include:

An attack on the host machine could compromise all guest machines operating on that host.

A bottleneck or failure of any hardware component that is shared between multiple guests, such as a failure in a disk subsystem, could affect multiple virtual machines.

While administration is centralized, virtualization is a newer technology and requires new skills, so managing virtual servers could add complexity.

A compromise of a guest system is typically limited to that system only because each virtual machine is kept partitioned from other guest machines. System isolation, if configured, is an advantage of virtualization. Isolation is typically used for testing purposes and prevents unreliable applications from interfering with other systems. Virtual systems do not need to be isolated. They can be configured to have full network access to other virtual machines or other network devices.

An advantage of virtualization is reduced hardware costs.

q_virt_type2_secp8

Which type of hypervisor runs as an application on the host machine?

Answers:

Type 1

***Type 2**

Type 3

Type 4

Explanation:

A Type 2 hypervisor is known as a hosted hypervisor. It runs as an application on a conventional operating system.

A Type 1 hypervisor is like a thin operating system that directly interfaces with the computer hardware.

There are no Type 3 or Type 4 hypervisors.

q_virt_virtualize_secp8

Which of the following are advantages of virtualization? (Select two.)

Answers:

***Centralized administration**

***Easy migration of systems to different hardware**

Redundancy of hardware components for fault tolerance

Reduced utilization of hardware resources

Improved host-based attack detection

Explanation:

Virtualization allows a single physical machine (known as the host operating system) to run multiple virtual machines (known as guest operating systems). The virtual machines appear to be self-contained and autonomous systems. Advantages of virtualization include:

Server consolidation

The ability to migrate systems between different hardware

Centralized management of multiple systems

Increase utilization of hardware resources

Isolation of systems and applications

Disadvantages of virtualization include:

A compromise in the host system could affect multiple guest systems.

A failure in a shared hardware resource could affect multiple systems.

q_virt_workload_secp8

Which load balancing method distributes a workload across multiple computers?

Answers:

Resource pooling

Virtualization

Bottleneck

***Workload balancing**

Explanation:

Workload balancing distributes a workload (the total requests made by users and applications of a system) across multiple computers or a computer cluster to achieve optimal resource utilization, maximum throughput, minimal response time, and less overload.

Resource pooling creates shared logical pools of CPU and memory resources from many physical machines within the hypervisor. This guarantees a level of resources for specific virtual machines.

Virtualization refers to installing and running multiple operating systems concurrently on a single physical machine.

A bottleneck is an area (software, hardware component, etc.) where all traffic slows down.

10.2 Virtual Networking

As you study this section, answer the following questions:

How does a virtual network differ from a physical network?

What is a Virtual Private Network (VPN)?

What is a virtual machine?

What terms are associated with virtualization and what do they mean?

What is the Dynamic Host Configuration Protocol (DHCP)?

How can physical devices become virtual ones?

Who are some of the network virtualization service providers?

In this section, you will learn to:

Configure virtual network devices.

Create virtual switches.

Key terms for this section include the following:

Term	Definition

Virtual network	A computer network consisting of virtual and physical devices.
Virtual local area network (VLAN)	A virtual LAN running on top of a physical LAN.
Virtual private network (VPN)	A secure tunnel to another network that connects multiple remote end-points.
Virtual machine (VM)	A virtual computer that functions like a physical computer.
Virtual switch (vSwitch)	Software that facilitates the communication between virtual machines by checking data packets before moving them to a destination.
Virtual router (vRouter)	Software that replicates the functionality of a physical router.
Virtual firewall appliance (vFA)	Software that functions as a network firewall device. A virtual firewall appliance provides packet filtering and monitoring functions.
Virtual machine monitor (VMM)/hypervisor	Software, firmware, or hardware that creates and runs virtual machines.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
TestOut Security Pro	3.3 Implement virtualization 3.3.1 Create virtual machines 3.3.2 Create virtual switches
CompTIA Security+ SY0-701	3.1 Compare and contrast security implications of different architecture models. Architecture and infrastructure concepts Network infrastructure Virtualization 4.1 Given a scenario, apply common security techniques to computing resources.

	Sandboxing

10.2.1 Virtual Networking Overview (Lesson Video)

Transcript:

One of the benefits of virtualization is the fact that the network environment can be virtualized, as well as the hardware environment. Each virtual machine on a hypervisor can be assigned a virtual network adapter. The virtual network adapter can be connected to either a physical network segment or to a virtual network segment. A virtual network segment has all the components that you would find on a physical network segment. You have the virtual network adapters, you have virtual network switches, and so on.

Let's take a look at how a virtual network works. Over here we have a typical physical network switch; we have traditional user desktops all interconnected using the switch and standard UTP cable, but notice over here that we also have a hypervisor host. This hypervisor host is also connected to the physical network switch with a UTP cable just like any other physical host would be. However, within the hypervisor we have four virtual machines running. One option for configuring virtual networking for these virtual machines is to take their virtual network adapters and bridge them logically--not with actual physical cables--but logically to the physical network interface installed in the hypervisor host, which of course is connected to the physical switch.

By doing this, each of these virtual machines then appears on the physical network as if they were a physical host, just like the systems over here are. The virtual adapter in each one of these virtual machines has their own MAC address and their own IP address. Therefore, they appear to be just a regular host on the physical network. This isn't the only option you have when networking virtual machines. In this example, we've created a virtual switch within the hypervisor. This is not a physical switch with real ports; instead, it's a switch that's running as software in the RAM of the hypervisor itself.

You can actually connect each one of these virtual network adapters and each one of these virtual machines to this virtual switch. We've depicted it here with what appears to be drop cables, but they're not really drop cables they're just logical electrical connections using the hardware, the hypervisor itself, to connect the virtual network adapters to this virtual switch. By doing this, we've created a virtual network segment.

Notice that this virtual switch is not connected in any way to the physical interface and the hypervisor at this time. Therefore, each one of these virtual machines can see each other on the virtual network through the virtual switch, but they cannot see any of these hosts out here on the physical network. This is sometimes called a sandbox environment. If you wanted to, you could actually connect the hypervisor itself to this virtual switch so the hypervisor can interact with these virtual machines on the virtual network. But it still keeps these virtual machines themselves isolated from the physical network and this can be a very valuable configuration, especially if you're in a testing environment.

For example, suppose you're about to roll out a new patch to all of your physical servers. You could set up virtual servers connected to an isolated virtual switch within the hypervisor and then test that update to see if it's going to break anything before you actually roll it out into your production environment. You can even create a router within your hypervisor.

In this case, we've set up a net router that connects our virtual network segment to the physical switch through the physical interface in the hypervisor. In this scenario, we've essentially created a virtual network segment that goes to the NAT router and then out on to your physical network segment. In this configuration, all of these internal virtual hosts are able to contact these physical hosts out here. However, because we're dealing with a NAT router, we can't initiate connections in the other direction of course, but you can go out in this direction. These hosts would be able to get out on the internet, for example, and download patches and updates.

There are a variety of different ways in which virtualization could be implemented. For example, you could implement virtualize servers. However, you could also virtualize desktops, using the virtual desktop interface, or VDI. There are actually many ways to do this. For example, you could create multiple virtual desktops on the same local workstation host. This is very useful in a situation where you're testing an application on multiple platforms to make sure that it runs properly. Another strategy would be to implement multiple desktops on one physical server, and then allow all of your end users to access those virtual desktops remotely.

Essentially, all of your user desktops are provided as virtual machines from a limited number of hypervisor hosts. This strategy could save you a lot of money in your desktop budget. You simply purchase a few high-end servers that will host the hypervisor and as desktop virtual machines. The end users on the other hand will require only very low-end workstations. In fact, they don't even need a hard drive if you have a basic operating system in the firmware of the system. These low-end workstations connect remotely to the hypervisor server and then run the desktop. You can even use virtualization to provide offsite virtualized network components.

For example, you could use virtualization to move your entire data center offsite. Instead of managing it onsite, it's taken care of by a vendor whose physical facility might be miles away from your physical location. In this scenario, the vendor manages your hardware and your software to provide the virtual network that's accessed by your end users. This can be very useful because it only requires very minimal hardware to be maintained at your physical site.

That's it for this lesson. It's important that you're familiar with how virtualization works. Almost all organizations are moving in that direction because of the cost savings associated with virtualization. In this lesson, we talked about how virtual bridging works. We talked about creating a sandbox with virtual networking. We talked about using a NAT router with virtual networking, and then we talked about virtualizing the network infrastructure.

10.2.2 Virtual Network Devices (Lesson Video)

Transcript:

Let's talk a little bit about virtual networks. A virtual network is a computer network consisting of virtual and physical devices. Virtual devices save organizations money. By using less physical storage space, a company is able to have at least twice as many devices in a network because they pay for very little space in a data center.

One form of network virtualization is protocol-based. You should already be familiar with this type of virtualization. VLANs and VPNs are two examples of protocol-based network virtualization.

Remember, with a VLAN, the network devices aren't physically segmented. The only thing that segments these devices is software defining how communications are routed by looking at the VLAN IDs of packets. The same is true with a VPN. There's no dedicated, physical connection between two devices using a VPN. It's a virtualized dedicated connection.

Now, let's go over some virtual networking devices and how they can be used to create a more secure network design.

The first network device we'll look at is a virtual switch, or vSwitch. A vSwitch is a software application that facilitates the communication between VMs by checking data packets before moving them to a destination. vSwitches are often integrated with virtual machine software. Sometimes, they're even part of a server's firmware.

The nice thing about a vSwitch is it's much easier to implement and manage than a traditional switch. In addition, a vSwitch can actually ensure the security integrity of virtual hosts, thereby creating a more secure network. For example, a vSwitch could make sure that a VM meets certain security criteria before it is allowed to communicate on the virtual network. If it fails the security check, its communications will be blocked.

Some of the most common vSwitch platforms are Open vSwitch (OVS), VMware virtual switch, and the Cisco Nexus 1000v. All three can function within a software-defined network as an SDN controller, allowing a much more dynamic network environment.

A virtual router, or vRouter, is a software function that replicates a physical router.

The nice thing about a virtual router is that it doesn't need to rely on the IP routing functionality used by physical routers. This means you can move routing functions around a network freely, creating a more dynamic network environment.

In addition, when you use a virtual router, you are free from a specific vendor ecosystem. This means you aren't locked in to proprietary applications or protocols and can more customize your networking environment. Virtual routers even work well with SDN infrastructure.

Note that, while virtual networking devices can be very flexible and offer a lot of benefits, they still aren't a replacement for physical networking devices.

Virtual networking devices use shared hardware resources, making them a poor choice for resource-intensive tasks, such as encrypting and decrypting information. Load balancing can also become an issue as the resources used to route, switch, and load-balance communication are shared among the devices.

Another drawback is all of these components rely on software applications and are susceptible to things like VM escape. If your entire network is virtualized, there is no physical boundary between certain systems. With physically segmented systems, it is impossible for an attacker to leap between segments without physical access. With virtual segmentation, an attacker only needs to know of an exploit that allows the leap.

That's it for this lesson. In this lesson, we looked at virtual networking devices. We looked at protocol-based virtualization, virtual switches, and virtual routers. We looked at how these virtual devices can benefit a networking environment. And we finished by looking at some of the drawbacks of virtual networking devices.

10.2.3 Configuring Virtual Network Devices (Demo Video)

Transcript:

In this demonstration, we're going to take a look at how to manage virtual devices. There are two types of virtual devices that we can easily manage from here.

One of them would be virtual switches, and I'm using the server version of Hyper-V, which is running on a Windows Server 2022. It's essentially the same as the client version, but all virtualization programs are going to have some type of virtual switch management. In Hyper-V, we're going to access the Virtual Switch Manager. The virtual switch should be created before you create your virtual machines. Now, in Hyper-V, we have different types of virtual switches. For example, I can create a private switch, and I'll just click 'Create Virtual Switch.' The private switch allows the virtual machines to communicate with each other, and that's it. If I wanted to establish an internal network, they could communicate with other virtual machines, interact among themselves and with the host, but they can't access the external network.
Alternatively, I could set up an external network where they can actually use the host's network card directly. You can see there's a check mark here for 'Allow management operating system to share this network adapter.' This means I'm using the same network adapter for both the virtual machines and the host. In a real server environment, it would be best practice to install a couple of network cards, allowing the host to use one network card and the virtual machines to use the other network card.
I can also enable VLANs on my virtual switches if I'm using VLANs. I've created a couple of private switches here, 'private' and 'private 2.' I also have an external switch called 'Internet.' I have plenty to work with. After setting up your virtual switches, the next step is to connect the virtual network cards within the virtual machines to the virtual switch. We're going to examine the settings in this test VM that I have here to see what happens.
I already have a synthetic network adapter installed in this machine, and it's connected to the private network. Again, if I'm using VLANs, I can enable VLAN identification. Depending on your virtualization system, you may have other additional features.
For instance, I can enable bandwidth management to prevent this virtual machine from monopolizing the network card's bandwidth. There are also additional features available when expanding the network card settings. For example, I can configure IPsec offloading, allowing the network card to handle some of the IPsec work when encryption using IPsec is in use. I can specify the details of how it should be handled. Furthermore, there are advanced features where I can customize the MAC address or enable MAC spoofing, which is required for supporting network load balancing. I also have DHCP guard to prevent this virtual machine from acting as a rogue DHCP server and router guard to stop it from becoming a rogue router. Additionally, I can designate it as a protected network, meaning that if its network is disconnected and this machine is part of a high availability solution, it would fail over to the other node in that cluster. Port mirroring allows me to send a copy of network traffic to another machine, which can be useful for monitoring purposes. For instance, if I'm using a product like Network Monitor or Wireshark. I can also implement NIC Teaming, allowing me to use multiple adapters in the guest operating system and link them together for either increased bandwidth or fault tolerance.
In this demonstration, we examined the various options for managing virtual devices.

10.2.4 Virtualization Implementation Facts

This lesson covers the following topics:

- Virtual networking
- Networking virtualization providers

Virtual Networking

A virtual network is made up of one or more virtual machines configured to access local or external network resources. Important facts about virtual networks include the following:

- Virtual machines support an unlimited number of virtual networks. Also, be aware that an unlimited number of virtual machines can be connected to a virtual network.
  - Multiple virtual networks can be associated with a single physical network adapter.
  - When a virtual network is created, its configuration is dependent on the configuration and physical hardware (such as the type and number of network adapters) of the host operating system.
  - The physical devices are partitioned into one or more virtual devices, depending on the network necessity and the device's capability.
  - When setting up a new virtual device, the system administrator will define how much of the physical device capability each partition will have. This means that one physical server could act as two or three virtual machines that work separately from one another and have their own specifications.
  - The available resources in a network are split up so the available bandwidth is turned into channels. Each channel can be assigned to a particular server or device in real-time. Each channel is independently secured.
- A virtual network includes a virtual Dynamic Host Configuration Protocol (DHCP) server that can provide IP address leases only to virtual machines. Even though the DHCP server is isolated, it assigns unique IP addresses from the range specified.
- Accessing a network and network resources requires that the operating system on the virtual machine be configured as a part of the network.
- Internal network virtualization configures a single system with software containers, or pseudo-interfaces, to emulate a physical network with software. This can improve a single system's efficiency by isolating applications to separate containers or pseudo-interfaces.
- External network virtualization combines one or more LANs into virtual networks to improve a large network's efficiency. Using this technology, systems physically attached to the same local network can be configured to be separate virtual networks. Systems from separate LANs can also be combined into a single VLAN that spans segments of a large network.
- Network virtualization should allow a virtual network, including all of its IP addresses, routes, network appliances, and so on, to appear to be running directly on the physical network. This allows the servers connected to that virtual network to continue to operate as if they were running directly on the physical network, even though multiple virtual networks share the physical network.

Network Virtualization Providers

Some of the main network virtualization service providers are:

Provider	Description
VMware	Be aware of the following regarding VMware solutions. VMware introduced the first x86 server virtualization products in 2001, making it a virtualization pioneer. VMware desktop software runs on Microsoft Windows, Linux, and macOS. In contrast, its enterprise software hypervisor for servers, VMware ESXi, is a bare-metal hypervisor that runs directly on server hardware without requiring an additional underlying operating system. ESXi is primarily used for data center virtualization.
Microsoft	Microsoft solutions include: Hyper-V Network Virtualization that provides virtual networks to virtual machines. This is similar to how server virtualization (hypervisor) provides virtual machines to the operating system. Hyper-V Network Virtualization has high scalability, with a capacity for over 1,000 virtual machines per host. Microsoft Azure that provides network virtualization in the cloud.
Citrix	Citrix virtualization solutions: Provide a virtualization solution called XenServer, also referred to as Citrix Hypervisor. Support the widest range of graphics applications. Support Intel GVT-g GPU virtualization, a CPU-embedded GPU requiring no additional hardware.

10.2.5 Virtual Networking Facts

This lesson covers the following topics:

- Virtual networks
- Virtual networking devices

Virtual Networks

A virtual network is a computer network consisting of virtual and physical devices. Organizations generally use virtual devices to save money. By using less physical storage space, a company is able to have considerably more devices in a network while using very little space in a data center. With virtualization, companies can take advantage of the efficiencies and agility of software-based devices and storage resources.

The physical networking devices are responsible for the forwarding of packets, while the virtual network (software) provides an intelligent abstraction that makes it easy to deploy and manage network services and underlying network resources.

The following are some network virtualization terms to be familiar with:

Term	Description
Virtual local area network (VLAN)	Several physical LANs can function as a single logical LAN or the partitioned network can be on a single router.
Virtual area network (VAN)	This is a virtual LAN running on top of a physical LAN. This configuration enables guest virtual machines on separate physical hosts to communicate.
Virtual private network (VPN)	A VPN is usually used as a secure tunnel over another network, connecting multiple remote end-points, such as routers. A multipoint VPN is a VPN connecting more than two end-points.
Virtual machine (VM)	VMs are virtual computers that function like a physical computer. Virtual servers are virtual machines capable of providing services such as databases, email, domains, and applications. The traffic between virtual machines can be routed using virtual switches alongside virtual routers and virtual firewalls for network segmentation and data isolation.

Virtual Networking Devices

The following table describes virtual networking devices that can be used to create a more secure network.

Device	Description
Virtual switch (vSwitch)	Software that facilitates the communication between virtual machines by checking data packets before moving them to a destination. A vSwitch may be a part of software installed in the virtual machine or part of the server firmware.
Virtual router (vRouter)	A software function that replicates the functionality of a physical router. Because virtual routing liberates the IP routing function from specific hardware, you can more freely move routing functions around a network.
Virtual firewall appliance (VFA)	Software that functions as a network firewall device that provides the usual packet filtering and monitoring. The VF can run as a traditional software firewall on a virtual machine.
Virtual machine monitor/hypervisor (VMM/hypervisor)	Software, firmware, or hardware that creates and runs virtual machines. A computer on which a hypervisor runs to provide one or more virtual machines is called a host machine. Each virtual machine is called a guest machine. The hypervisor provides the guest operating systems with a virtual operating platform and manages the execution of the guest operating systems.

10.2.6 Create Virtual Switches (Simulation)

Scenario

You have installed Hyper-V on the CorpServer server. You want to use the server to create virtual machines. Prior to creating the virtual machines, you are experimenting with virtual switches.

In this lab, your task is to:

Create an internal virtual switch named *Switch 1.*

Create a private virtual switch named *Switch 2.*

Explanation

Complete this lab as follows:

Open the Virtual Switch Manager.

From Hyper-V Manager, right-click **CORPSERVER** .

Select **Virtual Switch Manager** .

Create an internal switch named *Switch 1* .

Select **Create Virtual Switch** .

In the Name field, enter **Switch 1** .

Under Connection type, select **Internal network** .

Select **Apply** .

Create a private switch named *Switch 2* .

From the left pane, select **New virtual network switch** .

From the right pane, select **Private** .

Select **Create Virtual Switch** .

In the Name field, enter **Switch 2** .

Select **OK** .

10.2.7 Practice Questions (Section Quiz)

q_virt_impl_citrix_secp8

Which of the following provides the network virtualization solution called XenServer?

Answers:

Microsoft

***Citrix**

VMWare

Cisco

Explanation:

Citrix provides a virtualization solution called XenServer, also referred to as Citrix Hypervisor.

Microsoft provides a virtualization solution called Hyper-V Network Virtualization.

VMWare provides a virtualization solution called ESXi.

Cisco does not provide a virtualization solution but does offer a vSwitch platform called Nexus 1000v.

q_virt_impl_hyper-v_secp8

You are a system administrator for a large corporation that is planning to implement network virtualization across its data centers.

The corporation has a high demand for scalability, with the need to support over 1,000 virtual machines per host. The corporation also has a significant investment in Microsoft technologies and wants to leverage existing resources and expertise.

Which network virtualization provider should you recommend?

Answers:

VMware

***Hyper-V**

Citrix

VLAN

Explanation:

Hyper-V Network Virtualization is the correct answer. It provides virtual networks to virtual machines, similar to how server virtualization provides virtual machines to the operating system. It has high scalability, with a capacity for over 1,000 virtual machines per host, which matches the corporation's needs. Additionally, as a Microsoft solution, it would allow the corporation to leverage its existing investment in Microsoft technologies and expertise.

While VMware is a pioneer in virtualization and offers robust solutions, it may not be the best fit for a corporation with a significant investment in Microsoft technologies. VMware's ESXi is primarily used for data center virtualization, but the question does not provide information about the need for this specific feature.

Citrix is a company that sells virtualization software that enables users to work from remote locations. However, it uses Xen server as its hypervisor, which makes it difficult to use with Microsoft technologies, including existing resources and expertise.

VLAN, or Virtual Local Area Network, is a network topology configuration that allows for the partitioning and isolation of devices in a larger network. While it can be used in conjunction with network virtualization, it is not a network virtualization provider itself.

q_virt_impl_micro_secp8

Which of the following is a network virtualization solution provided by Microsoft?

Answers:

VMware

VirtualBox

***Hyper-V**

Citrix

Explanation:

Hyper-V Network Virtualization provides virtual networks to virtual machines. This is similar to how server virtualization (hypervisors) provides virtual machines to the operating system. Hyper-V Network Virtualization has high scalability, with the capacity for over 1,000 virtual machines per host.

None of the other virtualization solutions are provided by Microsoft.

q_virt_impl_virtual_networking_01_secp8

You are a system administrator for a large corporation that is planning to implement virtual networking to improve efficiency. You have been tasked with setting up a new virtual device.

Which of the following steps would be the most appropriate to ensure optimal utilization of the physical device capability?

Answers:

Assign all available resources of the physical device to the new virtual device.

***Partition the physical device into multiple virtual devices, each with a portion of the physical device's capability.**

Assign a small portion of the physical device's resources to the new virtual device to ensure other devices can also use the resources.

Avoid partitioning the physical device and use it as a single virtual device.

Explanation:

Partitioning the physical device into multiple virtual devices, each with a portion of the physical device's capability, is correct. This allows for optimal utilization of the physical device's resources as it ensures that each virtual device has the resources it needs without monopolizing the entire physical device.

Assigning all available resources of the physical device to the new virtual device is incorrect. This would not allow for optimal utilization as it would prevent other virtual devices from using the physical device's resources, potentially leading to resource wastage if the new virtual device does not require all the resources.

Assigning a small portion of the physical device's resources to the new virtual device to ensure other devices can also use the resources is incorrect. While it is important to ensure that resources are available for other devices, assigning a small portion may not provide the new virtual device with the resources it needs to function effectively.

Avoiding partitioning the physical device and using it as a single virtual device is incorrect. This would not allow for optimal utilization of the physical device's resources as it would prevent the creation of additional virtual devices, limiting the overall efficiency and flexibility of the network.

q_virt_impl_virtual_networking_02_secp8

As a network administrator for a tech startup, you are tasked with improving the efficiency of a single system that runs multiple applications.

The system is currently experiencing performance issues due to the applications competing for network resources.

Which type of network virtualization would be the most appropriate solution in this scenario?

Answers:

External network virtualization

Virtual Private Network (VPN)

Virtual Local Area Network (VLAN)

***Internal network virtualization**

Explanation:

Internal network virtualization is correct. This type of virtualization configures a single system with software containers, or pseudo-interfaces, to emulate a physical network with software. This can improve a single system's efficiency by isolating applications to separate containers or pseudo-interfaces, thus reducing competition for network resources.

External network virtualization is incorrect. This type of virtualization is used to combine one or more LANs into virtual networks to improve a large network's efficiency. It is not typically used for improving the efficiency of a single system running multiple applications.

Virtual Private Network (VPN) is incorrect. While a VPN can provide secure access to a private network from a remote location, it does not specifically address the issue of improving the efficiency of a single system running multiple applications.

Virtual Local Area Network (VLAN) is incorrect. A VLAN is used to group workstations that are not within the same geographical location into the same broadcast domain. It does not specifically address the issue of improving the efficiency of a single system running multiple applications.

q_virt_impl_vmware_esxi_secp8

Which of the following is a bare-metal hypervisor that runs directly on server hardware without requiring an additional underlying operating system and is primarily used for data center virtualization?

Answers:

VMware Workstation

***VMware ESXi**

VMware Fusion

VMware Horizon

Explanation:

VMware ESXi is correct because it is a bare-metal hypervisor that runs directly on server hardware without requiring an additional underlying operating system. It is primarily used for data center virtualization.

VMware Workstation is incorrect because it is a hosted hypervisor that runs on x64 versions of Windows and Linux operating systems, allowing users to set up virtual machines (VMs) on a single physical machine and use them simultaneously along with the actual machine.

VMware Fusion is incorrect because it is a software hypervisor developed by VMware for Macintosh computers with Intel processors, allowing users to run Windows, Linux, and other operating systems simultaneously with Mac OS X on Intel-based Macs.

VMware Horizon is incorrect because it is a product suite that provides virtual desktops and applications to end users. It does not run directly on server hardware without an underlying operating system.

q_virt_net_network_secp8

Which of the following statements about virtual networks is true? (Select two.)

Answers:

A virtual network is independent of the configuration and physical hardware of the host operating system.

***A virtual network is dependent on the configuration and physical hardware of the host operating system.**

Each virtual network must be associated with a single physical network adapter.

***Multiple virtual networks can be associated with a single physical network adapter.**

Accessing network resources requires that the operating system on the virtual machine be configured on an isolated network.

Explanation:

A virtual network is made up of one or more virtual machines configured to access local or external network resources. Some important facts about virtual networks include:

Virtual machines support an unlimited number of virtual networks, and an unlimited number of virtual machines can be connected to a virtual network.

Multiple virtual networks can be associated with a single physical network adapter.

When a virtual network is created, its configuration is dependent on the configuration and physical hardware (such as the type and number of network adapters) of the host operating system.

Accessing a network and network resources requires that the operating system on the virtual machine be configured as a part of the network.

q_virt_net_switch_secp8

Which of the following devices facilitates communication between different virtual machines by checking data packets before moving them to a destination?

Answers:

***Virtual switch**

Virtual router

Virtual firewall

Hypervisor

Explanation:

A virtual switch is software that facilitates the communication between different virtual machines. It does so by checking data packets before moving them to a destination. They may already be a part of software installed in the virtual machine, or they may be part of the server firmware.

A virtual router is a software function that replicates the functionality of a physical router and contributes to the functioning of a virtual switch. However, it does not focus on checking data packets before moving them to a destination.

A virtual firewall is software that functions as a network firewall device that provides the usual packet filtering and monitoring. It does not check data packets before moving them to a destination.

The hypervisor is software, firmware, or hardware that creates and runs virtual machines, including virtual switches.

q_virt_net_van_secp8

Which of the following is a virtual LAN that runs on top of a physical LAN?

Answers:

VLAN

***VAN**

VFA

VMM

Explanation:

A virtual area network (VAN) is a virtual LAN running on top of a physical LAN. This configuration enables guest virtual machines on separate physical hosts to communicate.

VLANs allow several physical LANs to function as a single logical LAN.

A VFA is a virtual firewall appliance. This is software that functions as a network firewall device.

A virtual machine monitor is software, firmware, or hardware that creates and runs virtual machines. This is also known as a hypervisor.

q_virt_net_vfa_secp8

Which of the following virtual devices provides packet filtering and monitoring?

Answers:

VLAN

vSwitch

***VFA**

VMM

Explanation:

A VFA is a virtual firewall appliance. This is software that functions as a network firewall device that provides the usual packet filtering and monitoring. A VFA can run as a traditional software firewall on a virtual machine.

VLANs allow several physical LANs to function as a single logical LAN.

A vSwitch is software that facilitates the communication between virtual machines by checking data packets before moving them to a destination.

A virtual machine monitor is software, firmware, or hardware that creates and runs virtual machines. This is also known as a hypervisor.

q_virt_net_vlan_secp8

Which of the following is an example of protocol-based network virtualization?

Answers:

***VLAN**

vSwitch

VFA

VMM

Explanation:

VLANs and VPNs are two examples of protocol-based network virtualization.

A vSwitch is software that facilitates the communication between virtual machines by checking data packets before moving them to a destination.

A VFA is a virtual firewall appliance. This is software that functions as a network firewall device.

A virtual machine monitor is software, firmware, or hardware that creates and runs virtual machines. This is also known as a hypervisor.

q_virt_net_vpn_secp8

Which of the following is used as a secure tunnel to connect two networks?

Answers:

VLAN

VAN

VFA

***VPN**

Explanation:

A virtual private network (VPN) is usually used as a secure tunnel over another network, connecting multiple remote endpoints (such as routers). A multipoint VPN is a VPN connecting more than two endpoints.

VLANs allow several physical LANs to function as a single logical LAN.

A virtual area network (VAN) is a virtual LAN running on top of a physical LAN.

A VFA is a virtual firewall appliance. This is software that functions as a network firewall device.

10.3 Software-Defined Networking

As you study this section, answer the following questions:

Which three layers exist in the software-defined networking (SDN) architecture?

What is the function of the controller?

What technology allows network and security professionals to manage, control, and make changes to a network?

What are the advantages of SDN?

What are the disadvantages of SDN?

Key terms for this section include the following:

Term	Definition
Software-defined networking	An architecture that allows network and security professionals to manage, control, and make changes to a network.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	3.1 Compare and contrast security implications of different architecture models. Architecture and infrastructure concepts Software-defined networking (SDN)

10.3.1 Software-Defined Networking Basics (Lesson Video)

Transcript:

Even though software-defined networking, also know as SDN, is a relatively new approach to networking, its use has grown exponentially the last few years.

The idea behind SDN is that network engineers should be able to use software to configure and intelligently control the network, rather than relying on the individual static configuration files that are located on each network device. By using SDN, network engineers are given the ability and flexibility to control their networks programmatically in order to accommodate rapidly changing environments.
For example, let's say you work in a datacenter. In this datacenter, you need a lot of equipment to create a network servers, switches, bridges, load balancers, cables, et cetera.
All of these devices are highly capable of making truly intelligent decisions within their individual capabilities. However, they require individual configurations on each of the devices. And if you need to make a change that affects multiple devices, then you need to manually make that change on each of these devices. In addition, making all of these devices work together perfectly can be a bit of a headache. For example, devices from different vendors might not work properly with each other.
With SDN, however, the configuration, management, and communication of all these devices is standardized and centralized. Software-defined networking takes all of these devices and combines them in much the same way a computer combines internal hardware components. It does this by using what's called a software-defined network controller.
An SDN controller functions a lot like traditional desktop operating system. It is able to inventory hardware components in the network, gather network statistics, make routing decisions based on gathered data, and facilitate communication between devices from different vendors. It can also be used to make wide-spread configuration changes on just one device. In essence, it's sort of like an operating system for the network.
For example, take our datacenter here. Let's say this datacenter suddenly gets an influx of user connections, and it needs to spin up new virtual machines to handle the load. Or perhaps it needs to offload some of these connections to another datacenter.
In a traditional networking environment, spinning up new VMs, routing connections to these VMs, and propagating these changes throughout the network could take too much time, and users wouldn't be able to use their resources.
With SDN, however, this influx of users is identified by the controller. The network administrator can then make configuration changes on the controller to redirect users to the new VMs. And because all these configurations are controlled at a central location the controller they are propagated out to the network almost instantaneously, allowing the

users to have no interruption in their connections. When the influx of user connections returns to normal, the controller can then return the network to the state it was in before.
One of the primary advantages to using an SDN controller is that it can also act autonomously in situations like this. For example, if this scenario occurs again, the controller can remember the actions that were taken and perform them automatically.
So, it's pretty easy to see why software-defined networking is here to stay.
That's it for this lesson. In this lesson, we learned the basics of software-defined networking, including how it works and how the controller communicates with network devices and responds to events like an influx of user connections.

10.3.2 SDN Infrastructure and Architecture (Lesson Video)

Transcript:

In this lesson, we're going to take a deeper look at software-defined networking, or SDN. We'll go over how software-defined networking can enhance your infrastructure to be more responsive, scalable, and secure.

In a previous lesson, we looked at how the SDN controller can respond to changing network conditions and make changes to data flow. Let's take a look at the underlying SDN architecture to see how this is done.
So, we have our controller right here. Remember, the controller itself is just a software platform that contains other applications it's the network's operating system.
The SDN architecture consists of three layers. In the middle, there's the control layer. This is where the controller resides.
Above that is the application layer. This is where various applications reside. One of these could be an app that creates a network map of all the devices in your network. Another one could be a load balancer that stops and starts VMs as resource use increases.
The application layer and these various apps communicate with the control layer through what's called the northbound interface. These are sometimes called northbound APIs.
Below the control layer is the physical layer. This is where the physical networking devices reside. Even though it's called the physical layer, it's also where the virtual networking devices reside. The physical layer communicates to the control layer through the southbound interface. The individual networking devices use southbound APIs to communicate with the control plane, and vice versa.
Understand that this control layer, the control plane, removes that control plane from the physical networking devices. In traditional networks, each of these devices would have an integrated control plane located on the device. However, SDN removes this from the devices and creates a single control plane.
There's one last thing to be aware of regarding software-defined networks. As with the implementation of any new feature or technology, the transition comes with some inherent challenges.
For example, a centralized controller introduces a new target for attackers. Imagine a third-party application that uses the northbound API to gain complete knowledge of your entire network. That could be very dangerous. Therefore, it's important to be aware of proper security techniques and how these interfaces can be compromised.
That's it for this lesson. In this lesson, we took a deeper look into software-defined networking, or networking programmability. We also talked about some of the security concerns you need to be aware of with SDN.

10.3.3 SDN Facts

This lesson covers the following topics:

- Software-defined networking (SDN)
- SDN advantages and disadvantages

Software-Defined Networking (SDN)

Software-defined networking (SDN) is a recent technology that allows network and security professionals to manage, control, and make changes to a network. The idea is that network engineers are able to use software to configure and intelligently control the network rather than relying on the individual static configuration files located on each network device.

SDN uses a controller to manage the devices. The controller is software that is able to inventory hardware components in the network, gather network statistics, make routing decisions based on gathered data, and facilitate communication between devices from different vendors. It can also be used to make widespread configuration changes on just one device.

The SDN architecture consists of three layers. The table below shows the functions:

SDN Layer	Function
Application layer	The Application layer communicates with the Control layer through what is called the northbound interface. These are sometimes called northbound APIs.
Control layer	The Control layer receives its requests from the Application layer and then provides the Physical layer with its configuration and instructions.
Physical layer	The Physical layer, also known as the Infrastructure layer, communicates to the Control layer through the southbound interface. The individual networking devices use southbound APIs to communicate with the control plane and vice versa. Even though this layer is called the Physical layer, it is where both physical and virtual network devices sit.

This architecture allows for fully automated deployment (or provisioning) of network links, appliances, and servers. This makes SDN an important part of the latest automation and orchestration technologies.

SDN Advantages and Disadvantages

Some advantages of SDN include:

- Centralized management
- More granular control
- Lower overall cost and labor
- Give new life to old networking hardware
- Gather network information and statistics
- Facilitate communication between hardware from different vendors

Some disadvantages of SDN include:

- It is currently a new technology
- Lack of vendor support
- Standards are still being developed

Centralized control opens a new target for security threats

10.3.4 Practice Questions (Section Quiz)

q_sdn_advantage_secp8

Which of the following is an advantage of Software Defined Networking (SDN)?

Answers:

It's currently a new technology.

Standards are still being developed.

***More granular control.**

It lacks vendor support.

Explanation:

One of the advantages of SDN is more granular control.

Some disadvantages of SDN include:

It's currently a new technology.

It lacks vendor support.

Standards are still being developed.

q_sdn_app_01_secp8

Which of the following BEST describes the Application SDN layer?

Answers:

Communicates with the Control layer through the Southbound Interface.

***Communicates with the Control layer through the Northbound Interface.**

Receives requests and then provides configuration and instructions.

Is software that is able to inventory hardware components in the network.

Explanation:

The Application layer communicates with the Control layer through what is called the Northbound Interface. These are sometimes called Northbound APIs.

The Physical layer, also known as the Infrastructure layer, communicates with the Control layer through the Southbound Interface.

The Control layer receives requests from the Application layer and then provides the Physical layer with its configuration and instructions.

The controller is software that is able to inventory hardware components in the network.

q_sdn_app_02_secp8

Which SDN layer would a load balancer that stops and starts VMs as resource use increases reside on?

Answers:

Physical

***Application**

Session

Control

Explanation:

Applications reside on the Application layer. A load balancer that stops and starts VMs as resource use increases is an example of an application that would reside on this layer.

The Physical layer is where both physical and virtual network devices sit.

The Session layer is the fifth layer of the OSI model.

The Control layer is the middle layer. This is where the controller resides.

q_sdn_controller_secp8

Software Defined Networking (SDN) uses a controller to manage devices. The controller is able to inventory hardware components on the network, gather network statistics, make routing decisions based on gathered data, and facilitate communication between devices from different vendors. It can also be used to make widespread configuration changes on just one device.

Which of the following BEST describes an SDN controller?

Answers:

***The SDN controller is software.**

The SDN controller is hardware.

The SDN controller is a virtual networking device.

The SDN controller is a networking protocol.

Explanation:

SDN uses a controller to manage devices. The controller is software that is able to inventory hardware components on the network, gather network statistics, make routing decisions based on gathered data, and facilitate communication between devices from different vendors. The controller can also be used to make widespread configuration changes on just one device.

SDN is not a virtual networking device, a networking protocol, or hardware.

q_sdn_def_secp8

A company is deploying a software service to monitor traffic and enforce security policies in its cloud environment.

Considering the need for responsiveness, which technology should the company consider using?

Answers:

***Serverless platforms and Software Defined Networking (SDN)**

Infrastructure as code (IaC)

Microservices

Virtual private cloud (VPC)

Explanation:

Serverless platforms allow the company to focus on application development by managing the underlying infrastructure, while SDN helps in efficient traffic monitoring and security policy enforcement.

While IaC is ideal for automated and consistent deployment and management of infrastructure, it does not inherently provide traffic monitoring or security policy enforcement. Furthermore, IaC may not ensure system responsiveness.

Although microservices support modular and independent application development, they do not directly address traffic monitoring, security policy enforcement, or system responsiveness requirements.

Traditional VPC offerings provide some level of security in a cloud environment but may lack the flexibility, scalability, and efficiency of serverless platforms. Additionally, they do not inherently offer effective traffic monitoring or system responsiveness.

q_sdn_layer_secp8

Drag the software defined networking (SDN) layer on the left to the appropriate function on the right. (Each SDN layer may be used once, more than once, or not at all.)

Answers:

Application layer

Control layer

Physical layer

Explanation:

The SDN architecture consists of three layers:

Application layer - communicates with the Control layer through the northbound interface. These are sometimes called northbound APIs.

Control layer - receives its requests from the Application layer and then provides the Physical layer with its configuration and instructions.

Physical layer - communicates with the Control layer through the. southbound interface. The individual networking devices use southbound APIs to communicate with the control plane and vice versa. Even though this is called the Physical layer, it is where both physical and virtual network devices sit. It is also known as the Infrastructure layer.

q_sdn_north_secp8

Which of the following does the Application layer use to communicate with the Control layer?

Answers:

***Northbound APIs**

Controllers

Southbound APIs

These layers do not communicate

Explanation:

The Application layer communicates with the Control layer through what is called the Northbound Interface. These are sometimes called Northbound APIs.

The controller is just a software platform that contains other applications. It can be thought of as the network's operating system.

The individual networking devices on the Physical layer use Southbound APIs to communicate with the control plane and vice versa.

The Application and Control layers do communicate.

q_sdn_phys_secp8

Which of the following BEST describes the Physical SDN layer?

Answers:

***Also known as the Infrastructure layer.**

Sometimes called Northbound APIs.

Receives its requests from the Application layer.

Gives new life to old networking hardware.

Explanation:

The Physical layer is also known as the Infrastructure layer.

The Application layer is sometimes called a Northbound API.

The Control layer receives its requests from the Application layer.

One advantage of SDN is that it gives new life to old networking hardware.

q_sdn_software_secp8

Network engineers have the option of using software to configure and control the network rather than relying on individual static configuration files that are located on each network device.

Which of the following is a relatively new technology that allows network and security professionals to use software to manage, control, and make changes to a network?

Answers:

***Software Defined Networking (SDN)**

Control layer networking

Load balancing software

Infrastructure software networking

Explanation:

Software Defined Networking (SDN) is a relatively new technology that allows network and security professionals to manage, control, and make changes to a network. Network engineers are able to use software to configure and control the network rather than relying on individual static configuration files that are located on each network device.

The Control layer is one of three layers that comprise Software Defined Networking. The other layers are the Application layer and the Physical layer.

Load balancers can be a component of the Application layer.

The Physical layer can also be referred to as the Infrastructure layer.

q_sdn_south_secp8

Which APIs do individual networking devices use to communicate with the control plane from the Physical layer?

Answers:

Northbound

Northbound and Southbound

***Southbound**

None

Explanation:

Individual networking devices on the Physical layer use Southbound APIs to communicate with the control plane and vice versa.

The Application layer communicates with the Control layer through what is called the Northbound Interface.

q_sdn_standards_secp8

Which of the following is a disadvantage of Software Defined Networking (SDN)?

Answers:

***SDN standards are still being developed.**

SDN creates centralized management.

SDN gathers network information and statistics.

SDN facilitates communication between hardware from different vendors.

Explanation:

Some of the disadvantages of SDN include:

Still a newer technology

Lack of vendor support

Standards are still being developed

Centralized control opens a new target for security threats

Some of the advantages of SDN include:

Centralized management

More granular control

Lower overall cost and labor

Gives new life to old networking hardware

Gathers network information and statistics

Facilitates communication between hardware from different vendors

10.4 Cloud Services

As you study this section, answer the following questions:

What is the difference between a hybrid cloud and a community cloud?

What is the difference between infrastructure as a service (IaaS) and platform as a service (PaaS)?

Which two implementations are available for software as a service (SaaS)?

What services does cloud computing provide?

Which cloud computing model allows the client to run software without purchasing servers, data center space, or network equipment?

Key terms for this section include the following:

Term	Definition
Cloud	A metaphor for the internet.
Cloud computing	Software, data access, computation, and storage services provided to clients through the internet.
Public cloud	A cloud that is deployed for shared use by multiple independent tenants.
Private cloud	A cloud that is deployed for use by a single entity.
Community cloud	Platforms, applications, storage, or other resources that are shared by several organizations.
Hybrid cloud	A cloud deployment that uses both private and public elements.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.2 Summarize fundamental security concepts. Authentication, Authorization, and Accounting (AAA) Authenticating people 2.2 Explain common threat vectors and attack surfaces. Supply chain Managed service providers (MSPs) Vendors

	Suppliers 2.3 Explain various types of vulnerabilities. Cloud-specific Supply chain Service provider 3.1 Compare and contrast security implications of different architecture models. Architecture and infrastructure concepts Cloud Hybrid considerations Third-party vendors Serverless Virtualization Considerations 4.1 Given a scenario, apply common security techniques to computing resources. Hardening targets Cloud infrastructure 5.3 Explain the processes associated with third-party risk assessment and management. Vendor assessment Penetration testing Right-to-audit clause Independent assessments Vendor monitoring

10.4.1 Cloud Services Introduction (Lesson Video)

Transcript:

More and more networks are shifting to a decentralized resource approach by utilizing various cloud services. Because of this, it's important that you have a solid understanding of cloud concepts and how cloud services impact a network.

Chances are you've heard about the cloud. It's used in almost every conversation about networking or the internet. But what exactly is the cloud? Well, it's actually a lot of different things and a lot of them have been around for a while.
Let's say we have an enterprise network here. And within this network, we have several switches, routers, some servers, workstations, and a connection to the internet. Traditionally, each of these workstations would have all the necessary applications and hardware components in order to carry out daily tasks, and these servers would each have a single OS installed and provide a specific function, such as a file server or DHCP server.
However, the cloud changes how a traditional network model works. For example, let's say these workstations need to perform a lot of calculations. In a traditional networking model, each of these workstations would need the necessary software and hardware to match to carry out these tasks. That can get expensive. What the cloud does is outsource these tasks to other machines, which are either located on the same network or provided by a third party via the Internet.
In essence, the cloud is computing resources that are offered or delivered as a service over a network. These computing resources include performing calculations or computations, the ability to store data, or the ability to network resources together. The cloud name itself comes from the cloud-shaped symbol used to denote a complex infrastructure in system and network diagrams.
In order for a service to be considered a true cloud service, it needs to have five characteristics, which have been determined by the National Institute of Standards and Technology.
First, it needs to be an on-demand self-service, meaning users must be able to dynamically obtain resources.
Second, it needs to be provided over a network, typically the Internet, and be available on multiple platforms, such as mobile phones, tablets, laptops, and workstations. This is called broad network access.
Third, it needs to have resource pooling. This means multiple users can use the same resources. These resources can be physical or virtual, but they need to be able to scale dynamically according to demand. For example, if a server suddenly gets a thousand connections, it needs the ability to provision additional computing power or offload clients to different VMs or servers.
Fourth, the service needs to have rapid elasticity. This means the services can seamlessly provision resources and scale resource usage back depending on demand.
And finally, the fifth characteristic is what's known as a measured service--that is, resources can be controlled and optimized automatically. In addition, this also means that the service provider can collect information about the service.
Now, cloud services can be categorized based on the way they are deployed, and there are a few different cloud deployment models you should know about. The first is a public cloud.
A public cloud is made available to the public by a service provider. These services are sometimes free, but often use a pay-as-you-go model.
With a private cloud, the service consists of an infrastructure provisioned and operated solely for a single organization. Private clouds are typically hosted at the organization. However, there are some private cloud providers that provide remote hosting.
Sometimes several organizations share infrastructure for a specific community with common concerns. This is called a community cloud.
And the final cloud deployment model is a hybrid cloud. Hybrid clouds are becoming increasingly common. Hybrid clouds are a combination of private, community, or public clouds for different aspects of the services that are provided. For instance, a hybrid cloud service can provide special load balancing between different cloud types if there is a spike in resource usage. Hybrid clouds also use both on-site and off-site resource hosting.
Now that you have an understanding of some of the basic concepts of the cloud service model, let's take a look at the different types of cloud services that are available to organizations.
The first cloud service we'll look at is called Infrastructure as a Service (IaaS).
I want to point out that all cloud services end with as a service. This is a quick way for you to identify a cloud service. They almost all use this initialism model, too.
Now, Infrastructure as a Service provides an organization with infrastructure components that is, computer hardware components as a service. And these hardware components can be anything from networking devices to storage and data processing resources.
For example, instead of having to install multiple file servers with redundant hard drives in your server room, you can pay a company for storage space on their servers. This eliminates the need to maintain and secure local file servers.
And let's say you need to process large amounts of data or compile and test applications. Instead of needing the hardware infrastructure on site, you can outsource this to an IaaS company, such as AWS.

Another nice thing about IaaS is that it can scale as needed. For example, let's say you need to add more computing resources. Without IaaS, you would have to go out and buy all that hardware, set it up, configure it, et cetera. However, if you are using IaaS, you just tell the cloud company that you need more resources, and they are immediately provided.
One thing to know about using IaaS is that it only provides the hardware resources. It's still your job to handle all the configuration and setup and interface with the cloud infrastructure.
The next cloud service is Platform as a Service (PaaS).
Platform as a Service provides all the underlying resources that IaaS provides; however, it goes a step further. In addition to the hardware resources, PaaS also provides the configuration, setup, and management of the cloud platform. For example, it handles installing operating systems, connecting virtual networks, and creating an entire environment that is ready for you to use immediately. One common example of PaaS is for software developers. A lot of companies that offer PaaS have packages that are specifically designed for developing, testing, and deploying software programs. Some PaaS cloud service providers also provide additional features, such software design systems, software development resources, testing, and deployment resources. Like IaaS, PaaS provides scalability and flexibility; it can grow or shrink with your organization's needs.
Next is Software as a Service (SaaS). Software as a Service is designed for your average end user. It provides end users with the applications that they need to perform their day-to-day work.
Instead of installing the applications locally, they use a web browser to access programs. Some SaaS providers charge a monthly or annual fee. But often, with SaaS, you pay for only the applications that you need. For example, if a user needs a word processing application for nine hours a day, that is how much you pay for. If another user needs the same application for only three hours a month, you pay much less.
SaaS also allows you to try new software before committing to it. You can try it out, paying only for the time you use it. If you really like it, you can pay for more use. If you don't like it, you can look for something else without losing much money.
SaaS comes in two forms, simple multi-tenancy and fine-grained multi-tenancy. With simple multi-tenancy, each cloud service customer has their own dedicated resources segregated from other customers' resources.
Fine-grained multi-tenancy offers the same level of segregation as simple multi-tenancy, but all of the resources used by all customers are pooled together and shared. Your data is segregated, so other customers can't access it. However, your computing resources, applications, and other cloud resources are shared with other customers.
The last cloud service model we'll look at is Security as a Service (SECaaS).
Now, you might be thinking, wait a second, I'm a security professional--why would I want to outsource my security services to the cloud? This isn't what SECaaS does.
In one way, SECaaS is similar to SaaS it provides a subscription model for applications and software. However, instead of being used for productivity on workstations, the applications and software are specific to organizational security.
For example, you might subscribe to a SECaaS system that provides a high-level anti-virus detection service that includes continuous anti-virus updates. Other services provided by SECaaS include DDoS protection, intrusion detection, and authentication services. One of the most common SECaaS providers is CloudFlare.
SECaaS can sometimes be much more cost effective for an organization than having to purchase all the necessary hardware equipment and personnel to properly protect a network from viruses, malware, and instruction. However, it is still a necessity to have an on-site security professional.
The cloud is an extremely useful service that more and more organizations implement. Because of this, it's important that you understand the underlying structure of the cloud and what it means to your organization's systems and infrastructure. In addition, by better understanding the cloud deployment models, you can better protect your organization.

10.4.2 Enhancing Cloud Performance (Lesson Video)

Transcript:

In this video, we're going to look at five different ways the cloud is being enhanced with technologies for performance: fog computing, edge computing, serverless architecture, services integration, and resource policies. Let's get started.

Fog computing, also known as fogging, is the cloud plus the Internet of Things.
Fog is a new distributed architecture that extends services offered in the cloud to edge devices. Fog nodes are made up of intelligence, computing power, and storage resources. This can be in the form of routers or smart gateways that communicate with the cloud.

These fog nodes might, for example, process signals coming in from sensors in a certain part of a city to measure CO2 emissions and then send alerts.

Edge devices are designed to accomplish specific tasks using apps. They come in all shapes and sizes, from smartwatches to autonomous cars to smart buildings. Edge devices can also be routers, switches, or firewalls. Fog computing uses decentralized local network architectures to speed up the analysis and retrieval of data near the source. With growing 5G Wireless networking and embedded artificial intelligence, we're about to have billions of apps and devices coming online and a tsunami of data that needs to be processed. Fog can help with that.

With fog, critical core functions are distributed. Computing, communications, control, storage, and decision making are all close to where the data originates locally. Smart devices and sensors are being developed so the components can instantly interact via wireless networks.

Fog promises to be a revolutionary new way of connecting, living, and working. Its benefits over cloud computing include lower latency, greater security, and enhanced capacity. Since fog nodes are physically closer to devices in the network, users can enjoy much lower latency with instant responses to their computing needs.

Fog provides greater security mainly because the data travels a much shorter distance in a distributed network. With fog, the storage capacity of the cloud is also enhanced. For example, a smart vehicle in trial right now is currently generating 35 gigabytes of data per hour. Secure fog computing has the power and reach to process this amount of data.

In addition, fog has the monitoring, detecting, and reporting capabilities to provide real-time incident response for security breaches, allowing it to quickly detect and isolate threats and minimize disruption of services.

While fog computing has many positive aspects, it also inherits a cornucopia of security issues from cloud computing: account hijacking, denial of service, access control issues, unsecure APIs, system and application vulnerabilities, shared technology issues, and much more.

Now let's discuss Edge computing. Edge computing is performed at or near the source of the data, not in the cloud. It's a component of fog computing where the architecture is distributed with decentralized processing power. We capture, store, process, and analyze the data near the client, where the data is usually generated.

From the wearable on your wrist to streaming video optimization, from controlling intersection traffic flow to smart homes and buildings, edge computing's future is exponentially expanding with the Internet of Everything, artificial intelligence, 5G, robotics, machine learning, and all the other technological developments taking place now.

For edge devices to be smart, they must look at incoming data, process it effectively, share it, and take some sort of action. Edge computing means the edge devices take these actions without requiring data to be transported elsewhere. Using a hub and spoke model, the cloud is the hub, and everything outside the spokes is the edge, like a spiderweb.

Think of it like this: edge computing is about processing the data locally, where cloud computing is focused on processing data in a remote data center that's a public or private cloud.

There are many benefits of edge computing. And much like fog computing, these include lower latency, increased bandwidth, greater resiliency, and data sovereignty.

Data sovereignty is the ability to own and manage your own data, and companies are very concerned about making sure their data remains theirs regardless of the technology they utilize.

Edge also empowers workloads that need processing in near real time, like autonomous cars communicating with each other. This speed and efficiency allows organizations to make more impactful and timely decisions, which reduces costs, helps them better engage with customers, and increases privacy.

Edge computing also suffers from many of the same security risks as fog and cloud computing.

Another way to enhancing cloud performance is by using serverless architecture, or serverless computing. It lets you build and run your apps and services without having to actually manage the infrastructure they run on. Yes, your apps are still running on a server, but you don't have to take care of it or its components.

Instead, your cloud provider takes care of all the server provisioning management, leaving your developers to focus on what they do best: creating their core product. And this saves them precious time and energy.

Basically, a serverless architecture eliminates most of the infrastructure management tasks of server provisioning, patching, OS and database maintenance, storage, and more. The cloud service provider takes care of fault tolerance and high availability.

Ultimately, serverless architecture reduces costs and frees the company and its developers to focus on product innovation and a faster time to market. This gives them great agility and responsiveness while they enjoy automatic application scaling and market competitiveness.

The serverless model is often referred to as Backend as a Service, or BaaS, meaning that companies simply pay for the resources consumed on the back end. Companies can also rent functions and compute runtimes, known as Function as a Service, or FaaS. This allows companies to execute application logic, but none of their data is stored. Major FaaS vendors are Amazon, Google, IBM, and Microsoft.

Benefits of cloud services integration include synchronized data and apps, increased agility, faster time to market, improved operational efficiency, reduced operational costs, increased flexibility, and scalability. Whether single-cloud or multi-cloud focused, businesses are enjoying the tremendous, feature-rich benefits of cloud services integration.
While there are many positive aspects to serverless, there are a few drawbacks.
The first is performance. Some code could suffer from latency because the cloud provider typically spins down the serverless code while it's not in use. If the code you need has to take time to start up, that will cause a delay. Also, serverless may not be appropriate for some workloads, like high-performance computing, due to resource limits.
The next drawback is security.
In serverless computing, the total attack surface is due to each component being an entry point to the serverless application. Some of these would include network sniffing, man-in-the-middle attack, phishing, SQL injection, DDoS, and code exploit. Further, because of limitations, companies can't implement an intrusion detection/prevention system, or IDS/IPS.
The next challenge is privacy. Many serverless environments use proprietary public cloud environments, and resources are often shared. To overcome this, serverless computing can be done on private clouds or even on-premeses, which allows companies full control over system privacy.
Integration with the cloud has really grown in popularity as Software as a Service solutions continue to increase. More businesses are running a hybrid mix of on-premises apps and SaaS apps, and this is creating a much greater need for innovative integration methods.
There are three major integration types of cloud technologies: cloud-to-cloud, cloud-to-on-prem, and a hybrid combination of the two. Data and app integration needs are the core concern when you're choosing from these strategies. Businesses can build their own integration solutions for their enterprise, or they can employ a third-party provider with a solution that's reusable, agile, and scalable. More than likely, integrating with the cloud could be far more cost-effective as well.
However, data and apps need to be protected from internal and external threats. Misconfigurations can open devastating security vulnerabilities. A good cloud access security broker, or CASB, can be a dedicated security solution that monitors and automates security configuration audits.
A resource policy is attached to a specific resource where you can specify who has access to it and what they can do with it. This helps to lock down applications and data on mobile devices and integrate with cloud resources. You can create policies in something like Microsoft Intune and then roll them out to all your users. This includes your users who would like to bring your own device, or BYOD.
You can keep your corporate apps and data secure and encrypted on these devices, regardless of the OS Android, IoS, or Windows. Policies are usually applied to groups of users or groups of devices. If you're using Microsoft Intune for your resource policies, you'll need to integrate with Azure Active Directory services in the Azure cloud.
That's it for this lesson. In this lesson, we looked at other technologies that enhance cloud performance like fog computing, edge computing, and serverless architecture. We also discussed services integration, and we ended by exploring resource policies. The cloud will continue to expand, and more options will become available to enhance its performance.

10.4.3 Cloud Computing Security Issues (Lesson Video)

Transcript:

Let's talk about cloud computing security issues. The benefits offered by cloud computingâ€”"the ability to use applications, platforms, and infrastructure as needed, and not use them when you don'tâ€”"is attractive. However, security is a key issue.

In a traditional system, you secure data behind layer after layer of security mechanismsâ€”"firewalls, ACLs, encryption, et cetera. And even after all that, it's still a challenge to keep your network safe. Add in the cloud component, and now all that data is being sent across the internet to a third-party organization who also needs to maintain their own system security.
Ultimately, the decision to utilize a cloud provider comes down to a partnership of trust. The cloud computing service provider has to ensure the security and privacy of your information. Security is part of the service they provide.
However, you, as a security administrator, are still responsible for protecting your company's data. One of the key things you need to remember is that industry and government regulations created to protect personal and business information still apply. Even though the data is being stored by the cloud computing provider, you're still responsible.

Whenever an organization is evaluating a move to the cloud, I always tell them that they need to approach the process of picking a cloud service provider in the same way they would buying a used car from a dealership. A used car may look great on the surface, but what you really need to be concerned with is what's under the hood. Take the same approach when selecting a cloud service provider.

Let's go over some of the key security considerations you need to keep in mind if you use a cloud service provider.

First, there's an issue of trust. You have to trust that your cloud service provider is good at information security. If your provider hasn't done a good job securing their environment, you could be in trouble. You can insist that the cloud services provider implement specific security measures, but you can't physically make them.

In addition, there's a monitoring issue. It's difficult to monitor security issues when your environment is provided by a cloud service provider. You're not on site, so you don't know what security issues they're facing. You don't see the midnight attack that they had to fight off last night. The cloud service provider probably doesn't want to let you know about it because they're afraid you might terminate their services.

Likewise, it's difficult to measure a cloud provider's degree of competence with information security because they probably don't want to expose their infrastructure to customers. Therefore, as we talked about with monitoring, it's hard to know how things are really going with your cloud computing environment. All you see really is the user interface.

Another thing to consider is authentication. Make sure that the cloud service provider authenticates everyone who's going to use the service. Don't choose the cloud service provider if they allow free access.

The next key thing to look for is partitioning. In other words, they make sure that your information and your data is kept completely separate from other customers' information and data. You don't want a fellow cloud computing customer gaining access to your data.

Within your organization, make sure that your individual users only have access to the applications that they need. Do not allow them to choose whatever they want to use.

One way to do this is to use a Cloud Access Security Broker (CASB). A CASB is a software tool or a service that sits between an organization and a cloud service provider. Its job is to make sure that all communication and access to the cloud service provider complies with the organization's security policies and procedures. For example, it might make sure devices connecting to the cloud meet certain security standards, such as full-device encryption, up-to-date malware protection, et cetera.

The next key thing to look for is updating. Make sure that the cloud service provider maintains a rigorous patch management program to ensure that the systems are updated. This will help protect your information from potential exploits.

After that, check out their processes. A good cloud computing service provider should have processes, policies, and procedures in place to handle various events. These processes, policies, and procedures should include everything, from something as simple as requesting access to a new application or something as critical as dealing with a major security breach. When you're evaluating a company, ask the company for a copy of these policies and procedures so you can see what they are going to do if a certain event happens.

In addition, verify that the cloud service provider has policies and procedures in place for monitoring security logs, monitoring unusual behavior, and performing the various tasks required to keep their system safe.

Next, make sure that they use some form of encryption to ensure that the data passing between your organization and their service is encrypted and can't be sniffed. Most likely, the data will pass through the untrusted network connection that is the internet. Secure data transmission is critical.

Last, the cloud computing service provider should probe their system on a regular basis. They should contract with a third party whose expertise is breaking into systems. The third party should probe the cloud computing service to see if there are any weaknesses that can be exploited.

Cloud services are here to stay. No matter the organization you work for, they are probably going to use some type of cloud service. This is because cloud services are very cost-effective and convenient. However, with convenience comes risk. And it's your job as a security professional to identify those risk, account for them, and minimize them as much as possible.

10.4.4 Cloud Computing Facts

This lesson covers the following topics:

- Cloud computing
- Types of clouds

Cloud computing models

Cloud security risk reduction

Threat actors and attack surfaces

Virtual Desktop Infrastructure (VDI)

Cloud Computing

Cloud computing is software, data access, computation, and storage services provided to clients through the internet. The term *cloud* is a metaphor for the internet. It is based on the basic cloud drawing used to represent the telephone network. It is now used to describe the internet infrastructure in computer network diagrams. Characteristics of cloud computing include:

Delivery of common business applications that are accessed from a web service or software (like a web browser).

The cloud connection can exist over the internet or a LAN.

Cloud computing does not require end-user knowledge of the physical location and configuration of the system that delivers the services.

Types of clouds

Cloud computing can be implemented in several different ways, including the following:

Type	Description
Public cloud	A *public cloud* can be accessed by anyone. Cloud-based computing resources, such as platforms, applications, storage, or other resources, are made available to the general public by a cloud service provider. The service provider may or may not require a fee for using these resources. For example, Google provides many publicly accessible cloud applications, such as Gmail and Google Docs.
Private cloud	A *private cloud* provides resources to a single organization. Access is restricted to the users within the organization. Private clouds can be hosted internally. Because of the expense and expertise required to implement, clouds are typically hosted externally by a third party. An organization commonly enters into an agreement with a cloud service provider, which provides secure access to cloud-based resources. The organization's data is kept separate and secure from any other organization using the same service provider.
Community cloud	A *community cloud* is designed to be shared by several organizations. Access is restricted to users within the organizations who are sharing the community cloud infrastructure. Community clouds can be hosted internally or on-premise, with each organization sharing the cost of implementation and maintenance. Because of the expense and expertise required, community clouds are commonly hosted externally by a third party.
Hybrid cloud	A *hybrid cloud* is composed of a combination of public, private, and community cloud resources from different service providers. The goal behind a hybrid cloud is to expand the functionality of a given cloud service by integrating it with other cloud services.

The advantages of cloud computing are:

- Flexible access.
- Ease of use.
- Self-service resource provisioning.
- API availability.
- Service metering.
- The ability to try software applications in cloud computing service models.

Cloud Computing Models

Cloud computing service models include the following:

Model	Description
Infrastructure as a Service (IaaS)	IaaS delivers infrastructure to the client, such as processing, storage, networks, and virtualized environments. The client deploys and runs software without purchasing servers, data center space, or network equipment.
Platform as a Service (PaaS)	PaaS delivers everything a developer needs to build an application. The deployment comes without the cost and complexity of buying and managing the underlying hardware and software layers.
Software as a Service (SaaS)	SaaS delivers software applications to the client over the internet or on a local area network. SaaS comes in two implementation types: Simple multi-tenancy in which each customer has its own resources that are segregated from other customers. Fine grain multi-tenancy segregates customers, but resources are shared.
Security as a Service (SECaaS)	SECaaS providers integrate their services into a corporate infrastructure. The applications and software are specific to organizational security. SECaaS is based on the Software-as-a-Service cloud computing model but is limited to information security services and does not require on-premises hardware. These security services can include authentication, anti-virus, anti-malware, spyware, intrusion detection, penetration testing, and security event management. SECaaS can sometimes be much more cost-effective for an organization than having to pay for all the necessary equipment and personnel to properly protect a network from viruses, malware, and instruction. However, it is still necessary to have an on-site security professional.

Cloud Security Risk Reduction

Cloud service providers reduce the risk of security breaches through the following actions.

- Authenticate all users who access the service and allow users to access only the applications and data that they need.
- Use a Cloud Access Security Broker (CASB). A CASB is a software tool or service that sits between an organization and a cloud service provider. Its job is to make sure that all communication and access to the cloud service provider complies with the organization's security policies and procedures.
- Segregate each organization's centrally stored data.
- Verify, test, and apply updates to the infrastructure.
- Establish a formal process for all facets of the service, from user requests to major data breaches and catastrophic events.
- Implement security monitoring for usage, unusual behavior, and other events.
- Implement encryption up to the point of use, such as the client's web browser.
- Probe for security holes with a third-party service provider.
- Comply with all regulatory measures, such as the Sarbanes-Oxley Act.

Threat Actors and Attack Surfaces

The attack surface is all the points at which a malicious threat actor could try to exploit a vulnerability. Any location or method where a threat actor can interact with a network port, app, computer, or user is part of a potential attack surface. Minimizing the attack surface means restricting access so that only a few known endpoints, protocols/ports, and services/methods are permitted. Each of these must be assessed for vulnerabilities and monitored for intrusions.

One potential threat to cloud services is what can be referred to as the supply chain. Because cloud services are used by organizations and other customers to create a number of possible end products and services, they can be considered a link in a supply chain. A supply chain is an end-to-end process of designing, manufacturing, and distributing goods and services to customers. Rather than attack the target directly, a threat actor may seek ways to infiltrate it via companies in its supply chain, such as an IaaS, PaaS, or SaaS.

The supply chain breadth and complexity expose organizations to a huge attack surface. For example, for a computer motherboard to be trustworthy, the supply chain of chip manufacturer, firmware code developer, OEM reseller, courier delivery company, and administrative staff responsible for provisioning the computing device to the end user must all be trustworthy. Anyone with the time and resources to modify the computer's firmware could create backdoor access. The same is true for any computer or network hardware, software, or service. This means that when employing companies offering cloud services, you must trust that their hardware, software, and services are secure.

Another important consideration regarding the supply chain is a Managed Services Provider (MSP). MSP provisions and supports IT resources such as networks, security, or web infrastructure. MSPs are useful when an organization finds it cheaper or more reliable to outsource all or part of IT provision rather than try to manage it directly. From a security point of view, this type of outsourcing is complex, as it can be difficult to monitor the MSP. The MSP's employees are all potential sources of insider threat.

Establishing a trusted supply chain for computer equipment and services essentially means denying malicious actors the time or resources to modify the assets supplied. For most businesses, using reputable vendors will represent the best practical effort at securing the supply chain. Government, military/security services, and large enterprises will exercise greater scrutiny. Particular care should be taken if using secondhand machines.

Virtual Desktop Infrastructure (VDI)

Cloud-based services can be hosted externally by third-party service providers or internally on your own virtualization infrastructure. For example, internal private clouds are commonly used to provide a VDI. Using VDI, user desktops are virtualized, running on high-end hardware in the data center instead of the end user's workstation hardware. The physical workstation is merely used to establish a remote connection to the user's virtualized desktop. This is sometimes called a *thin client* deployment because most of the computing power is provided by servers in the data center. Traditional deployments, where most of the processing load is handled by the local workstation, are called *thick client* deployments.

Using VDI provides increased flexibility, enhanced security, efficient management, and better data protection than the traditional workstation-based desktop model. Consider the following advantages:

- Workstation hardware costs are reduced. Only minimal workstation hardware is required to run a Remote Desktop (Windows) or VNC (Linux) client and connect to the private cloud.
- User data on the desktop can be protected centrally by backing up the hypervisors where the virtualized desktops are running. There is no need to back up physical workstations separately.
- If a user's physical workstation fails, no data is lost. The user can access the virtualized desktop from a different workstation while the failed hardware is repaired or replaced.
- If a widespread malware infection hits multiple user desktops, the affected virtual systems can be quickly re-imaged on the hypervisor. There is no need to push large images down to end users' workstations over the network.
- If a user loses a device, such as a notebook or tablet, there is much less of a chance that critical data will be compromised because no data is saved on the device.

10.4.5 Cloud Storage Security Facts

This lesson covers the following topics:

- Cloud storage
- Advantages of cloud storage

Cloud Storage

Cloud storage is a data storage model. It is usually provided by a third party as a service. Some of the most widely used cloud storage for enterprise providers are Google Cloud, Amazon Web Services, and Microsoft Azure. Many companies take advantage of cloud services to decrease costs and meet ever-increasing storage needs.

Cloud storage services may be accessed through a co-located cloud computer service, a web service application programming interface (API), or applications that utilize the API. Cloud desktop storage that uses a cloud storage gateway or web-based content management system is an example of an application that uses the API.

A cloud access security broker (CASB) may act as a gatekeeper, extending an organization's security policies into the cloud storage infrastructure. A CASB focuses on the visibility of the company data, regulation compliance, user access to prevent threats, and data security through encryption and loss prevention.

Cloud storage is:

- Made up of many distributed resources but still acts as one in a federated or cooperative storage cloud architecture.

Highly fault tolerant through redundancy and distribution of data.

Highly durable through the creation of versioned copies.

Cloud Storage is a virtual service; the infrastructure is the responsibility of the storage provider. Access controls should be set in the same way as a local file system would be set. There is no need for the provider to have access to the stored data. Measures for securing cloud storage include:

Implement security controls in the same way as in a physical datacenter.

Use data classification policies.

Assign information into categories that determine storage, handling, and access requirements.

Assign security classification based on information sensitivity and criticality.

Use specialized tools to dispose of data securely when it is no longer needed.

Advantages of Cloud Storage

The advantages of Cloud Storage are:

Companies pay only for the storage used. This does not necessarily mean that cloud storage is less expensive, but it incurs only operating expenses.

Cloud storage can cut energy consumption by up to 70%, making an organization more green.

Organizations can choose between off-premises and on-premises cloud storage options or a mixture of the two options.

Storage availability and data protection are intrinsic to object storage architecture. Depending on the application, you can eliminate the costs, effort, and additional technology to add availability and protection.

Storage maintenance tasks, such as purchasing additional storage capacity, are the responsibility of the service provider.

Cloud storage can be used to copy virtual machine images from the cloud to on-premises locations or import a virtual machine image from an on-premises location to the cloud image library.

Cloud storage can be used as natural disaster backup since cloud storage providers' backup servers are typically located in different places around the globe.

10.4.6 Analyze Infrastructure Types and Functions

10.4.7 Practice Questions (Section Quiz)

q_cloud_comp_characteristics_secp8

Which of the following is NOT a characteristic of cloud computing?

Answers:

Delivery of common business applications accessed from a web service or software.

The cloud connection can exist over the internet or a LAN.

***Cloud computing requires end-user knowledge of the physical location and configuration of the system that delivers the services.**

Cloud computing does not require end-user knowledge of the physical location and configuration of the system that delivers the services.

Explanation:

Cloud computing requires end-user knowledge of the physical location and configuration of the system that delivers the services is NOT a characteristic of cloud computing. One of the key features of cloud computing is that it abstracts the complexity and details of the system infrastructure from the end user. Users can access and use the services without needing to understand the underlying technology or infrastructure.

Delivery of common business applications accessed from a web service or software is a characteristic of cloud computing. Common business applications are delivered over the internet and can be accessed from a web service or software like a web browser.

The cloud connection can exist over the internet or a LAN is also a characteristic of cloud computing. The cloud connection can exist over the internet or a LAN (Local Area Network).

Cloud computing does not require end-user knowledge of the physical location and configuration of the system that delivers the services is a characteristic of cloud computing. Users do not need to know the physical location and configuration of the system that delivers the services.

q_cloud_comp_cloud_secp8

Match each description on the left with the appropriate cloud technology on the right.

Answers:

Provides cloud services to just about anyone.

Allows cloud services to be shared by several organizations.

Integrates one cloud service with other cloud services.

Provides cloud services to a single organization.

Explanation:

Cloud computing can be implemented in several different ways, including the following:

A public cloud can be accessed by anyone. Cloud-based computing resources are made available to the general public by a cloud service provider. The service provider may or may not require a fee for use of these resources. For example, Google provides many publicly accessible cloud applications, such as Gmail and Google Docs.

A private cloud provides resources to a single organization. Access is restricted to only the users within that organization. An organization commonly enters into an agreement with a cloud service provider, which provides

secure access to cloud-based resources. The organization's data is kept separate and secure from any other organization using the same service provider.

A community cloud is designed to be shared by several organizations. Access is restricted to only users within the organizations who are sharing the community cloud infrastructure. Community clouds are commonly hosted externally by a third party.

A hybrid cloud is composed of a combination of public, private, and community cloud resources from different service providers. The goal behind a hybrid cloud is to expand the functionality of a given cloud service by integrating it with other cloud services.

q_cloud_comp_community_secp8

A group of small local businesses have joined together to share access to a cloud-based payment system.

Which type of cloud is MOST likely being implemented?

Answers:

Hybrid

Public

Private

***Community**

Explanation:

A community cloud is designed to be shared by several organizations. Access is restricted to users within the organizations who are sharing the community cloud infrastructure.

A hybrid cloud is composed of a combination of public, private, and community cloud resources from different service providers.

A public cloud can be accessed by anyone.

A private cloud provides resources to a single organization.

q_cloud_comp_hybrid_secp8

A large organization is planning to move its operations to the cloud and is considering different cloud deployment models.

The organization wants to achieve a balance of cost, security, flexibility, and control over its data and applications and is considering a hybrid cloud model but has concerns about the security implications.

Which of the following is a potential security challenge the organization should consider when using a hybrid cloud model?

Answers:

***The organization may struggle with managing multiple cloud environments and enforcing consistent security policies.**

The organization will have no control over the security of its data and applications in the cloud.

The organization will have to bear the full cost of managing and securing the cloud infrastructure.

The hybrid cloud model is not scalable and may not be able to meet the organization's growing needs.

Explanation:

A hybrid cloud model can present security challenges, including the complexity of managing multiple cloud environments and enforcing consistent security policies across all environments.

In a hybrid cloud model, the organization can control the security of its data and applications in the private cloud infrastructure. However, it needs to ensure that the public cloud provider also has robust security measures.

In a hybrid cloud model, the organization would typically be responsible for securing the private cloud infrastructure, while the public cloud provider would be responsible for securing the public cloud infrastructure.

A hybrid cloud model allows the organization to use the private cloud for sensitive data and applications and use the public cloud for less sensitive, scalable workloads.

q_cloud_comp_iaas_secp8

A company is considering moving its applications and data to the cloud. The company handles sensitive data and wants to maintain control over the security of its applications and data.

It is considering using an infrastructure as a service (IaaS) model.

Which of the following is a key responsibility the company will need to manage in an IaaS model?

Answers:

Securing foundational elements of networking, such as DDoS protection

Physical security of the cloud infrastructure

***Protection of operating systems when deployed**

Cloud storage backup and recovery

Explanation:

In an IaaS model, the customer is responsible for protecting the operating systems it deploys on the cloud infrastructure. This includes tasks like applying security updates and patches, managing access controls, and implementing intrusion detection systems.

In an IaaS model, the cloud service provider is typically responsible for securing foundational elements of networking, including distributed denial of service (DDoS) protection.

The cloud service provider, not the customer, typically manages the physical security of the cloud infrastructure.

While customers must manage their data and backups, the cloud service provider typically provides and manages the infrastructure for cloud storage backup and recovery.

q_cloud_comp_paas_secp8

Which of the following BEST describes the platform as a service (PaaS) cloud computing service model?

Answers:

***PaaS delivers everything a developer needs to build an application on the cloud infrastructure.**

PaaS delivers infrastructure to the client, such as processing, storage, networks, and virtualized environments.

PaaS delivers software applications to the client either over the internet or on a local area network (LAN).

PaaS stores and provides data from a centralized location without the need for local collection and storage.

Explanation:

Platform as a service (PaaS) delivers everything a developer needs to build an application on the cloud infrastructure. The deployment comes without the cost and complexity of buying and managing the underlying hardware and software layers.

Software as a service (SaaS) delivers software applications to the client either over the internet or on a local area network. Infrastructure as a service (IaaS) delivers infrastructure to the client, such as processing, storage, networks, and virtualized environments. The client deploys and runs software without purchasing servers, data center space, or network equipment. Data as a service (DaaS) stores and provides data from a centralized location without the need for local collection and storage.

q_cloud_comp_risk_reduction_secp8

You are the CISO of a large organization that is transitioning to cloud services. One of your main concerns is ensuring the security of your organization's data in the cloud.

Which of the following actions would be the LEAST effective in reducing cloud security risks?

Answers:

Implementing a Cloud Access Security Broker (CASB) to ensure all communication and access to the cloud service provider complies with the organization's security policies and procedures.

Segregating each organization's centrally-stored data.

***Regularly updating the organization's social media accounts with information about the transition to the cloud.**

Establishing a formal process for all facets of the service, from user requests to major data breaches and catastrophic events.

Explanation:

Regularly updating the organization's social media accounts with information about the transition to the cloud would be the LEAST effective measure. While transparency is important, regularly updating social media accounts with information about the transition to the cloud does not directly contribute to reducing cloud security risks. In fact, it could potentially provide threat actors with information they could use to target the organization.

A CASB is a software tool or service that sits between an organization and a cloud service provider to ensure all communication and access complies with the organization's security policies and procedures.

Segregating each organization's centrally-stored data can help prevent unauthorized access and data breaches.

Establishing a formal process for all facets of the service can help ensure that potential security issues are addressed promptly and effectively, reducing the risk of data breaches and other security incidents.

q_cloud_comp_saas_01_secp8

Which of the following cloud computing solutions delivers software applications to a client either over the internet or on a local area network?

Answers:

IaaS

PaaS

***SaaS**

DaaS

Explanation:

Software as a service (SaaS) delivers software applications to the client either over the internet or on a local area network (LAN).

Infrastructure as a service (IaaS) delivers infrastructure to the client, such as processing, storage, networks, and virtualized environments. The client deploys and runs software without purchasing servers, data center space, or network equipment.

Platform as a service (PaaS) delivers everything a developer needs to build an application on the cloud infrastructure. The deployment comes without the cost and complexity of buying and managing the underlying hardware and software layers.

Data as a service (DaaS) stores and provides data from a centralized location without the need for local collection and storage.

q_cloud_comp_saas_02_secp8

A rock band wishes to set up a system for communicating with their fans upon arrival at concerts and providing them with relevant hashtags for participation.

Which type of cloud service model would be MOST beneficial to recommend to the rock band?

Answers:

***Software as a service (SaaS)**

Platform as a service (PaaS)

Infrastructure as a service (IaaS)

Third-party vendor

Explanation:

The software as a service (SaaS) model would be the most appropriate recommendation, enabling the band to use preconfigured applications for sending text messages. It also manages fan engagement without extensive development or infrastructure setup.

Platform as a service (PaaS) provides resources and platforms for developing software applications. However, the band does not need to develop its software in this scenario.

Infrastructure as a service (IaaS) allows for provisioning IT resources such as servers, load balancers, and storage components. However, this is unnecessary as the band does not require infrastructure.

A third-party vendor refers to providers offering cloud services. While this type of vendor can be part of the solution, it is not a service model.

q_cloud_comp_secaas_secp8

You are the security administrator for your organization. You have implemented a cloud service to provide features such as authentication, anti-malware, intrusion detection, and penetration testing.

Which cloud service have you MOST likely implemented?

Answers:

IaaS

***SECaaS**

PaaS

SaaS

Explanation:

Security as a service (SECaaS) providers integrate their services into a corporate infrastructure. The applications and software are specific to organizational security. SECaaS is based on the software as a service (SaaS) cloud computing model. However, it is limited to information security services and does not require on-premises hardware. These security services can include authentication, antivirus, anti-malware, spyware, intrusion detection, penetration testing, and security event management.

IaaS delivers infrastructure to the client, such as processing, storage, networks, and virtualized environments.

PaaS delivers everything a developer needs to build an application.

SaaS delivers software applications to the client over the internet or on a local area network.

q_cloud_comp_supply_chain_secp8

A recent attack on a major retail chain resulted in stolen customer private information, including credit card information.

The report explained that a heating, ventilation, and air conditioning (HVAC) contractor copied the information to an external hard drive while servicing an air conditioner unit and later uploaded the data to a cloud storage resource.

A security engineer would classify this type of attack as which of the following?

Answers:

***Supply chain attack**

Cloud-based attack

Birthday attack

USB cable attack

Explanation:

A supply chain attack involves a threat actor seeking methods to infiltrate a company in its supply chain. An HVAC supplier is one example of using a maintenance service to gain access to sensitive areas like a data center.

A cloud-based attack involves a threat actor compromising one account with access to cloud resources to compromise other cloud assets further.

A birthday attack is a brute-force attack aimed at exploiting collisions in hash functions.

A Universal Serial Bus (USB) cable attack involves accessing unsuspecting users after they try to plug their devices into malicious USB cables or plugs, similar to card skimmers.

q_cloud_comp_thin_secp8

The IT manager has tasked you with installing new physical machines. These computer systems are barebone systems that simply establish a remote connection to the data center to run the user's virtualized desktop.

Which type of deployment model is being used?

Answers:

***Thin client**

IaaS

PaaS

Thick client

Explanation:

This type of deployment is often referred to as a thin client deployment. This deployment utilizes virtual desktop infrastructure (VDI) to virtualize a user's desktop. The client machine is essentially only used to connect to the high-end machines in the data center.

IaaS delivers infrastructure to the client, such as processing, storage, networks, and virtualized environments.

PaaS delivers everything a developer needs to build an application.

Traditional deployments, where most of the processing load is handled by the local workstation, are called thick client deployments.

q_cloud_comp_vdi_secp8

Which of the following are true concerning virtual desktop infrastructure (VDI)? (Select two.)

Answers:

***User desktop environments are centrally hosted on servers instead of on individual desktop systems.**

User desktop environments are provided by individual desktop systems instead of by remote servers.

***In the event of a widespread malware infection, the administrator can quickly reimage all user desktops on a few central servers.**

In the event of a widespread malware infection, the administrator can reimage user desktops by pushing an image out to each user desktop system over the network.

Roaming profiles must be configured to allow mobile users to keep their same desktop environment across systems.

Explanation:

Virtual desktop infrastructure (VDI) is a service that hosts user desktop environments on centralized servers. Users access their desktops from low-end systems over a network connection using a remote display protocol such as Remote Desktop or Virtual Network Computing (VNC). This allows users to access their desktop environment with their applications and data from any location and from any client device. Roaming profiles are not needed.

VDI provides administrators with a centralized client environment that is easier and more efficient to manage. For example, if a widespread malware infection hits multiple user desktops, the affected systems can be quickly reimaged on the VDI server. There is no need to push large images down to client systems over the network.

q_cloud_stor_available_secp8

Google Cloud, Amazon Web Services (AWS), and Microsoft Azure are some of the most widely used cloud storage solutions for enterprises.

Which of the following factors prompt most companies to take advantage of cloud storage? (Select two.)

Answers:

***Growing demand for storage**

***Need to bring costs down**

Need for a storage provider to manage access control

Need for platform as a service (PaaS) for developing applications

Need for software as a service (SaaS) for managing enterprise applications

Explanation:

Some of the most widely used cloud storage for enterprises are Google Cloud, Amazon Web Services, and Microsoft Azure. Because of the growing demand for storage and desire to bring costs down, many companies have been taking advantage of cloud storage.

The majority of companies prefer to manage access control to data in the cloud in order instead of turning over this critical responsibility to a third-party, such as a storage provider.

While PaaS and SaaS are types of cloud environments that provide tools and applications for developing applications or managing enterprise applications, most companies want to reserve access control to the data stored on these platforms in order to maintain a high level of security.

q_cloud_stor_casb_secp8

Which of the following cloud storage access services acts as a gatekeeper, extending an organization's security policies into the cloud storage infrastructure?

Answers:

***A cloud-access security broker**

A web service application programming interface

A co-located cloud computer service

A cloud storage gateway

Explanation:

A cloud-access security broker (CASB) may act as a gatekeeper, extending an organization's security policies into the cloud storage infrastructure. A CASB focuses on the visibility of company data, regulation compliance, user access, and data security through encryption and loss prevention.

Cloud storage services may be accessed through a co-located cloud computer service, a web service application programming interface (API), or by applications that utilize the API, such as cloud desktop storage (in other words, cloud storage gateways or web-based content management systems).

q_cloud_stor_casb_sol_secp8

A large multinational corporation is planning to move a significant portion of their data storage to the cloud.

As the company's lead cybersecurity analyst, you are tasked with ensuring the security of the company's data during and after the transition.

Which of the following would be the most effective solution to extend the company's security policies into the cloud storage infrastructure, provide visibility of the company data, ensure regulation compliance, prevent threats through user access control, and secure data through encryption and loss prevention?

Answers:

Implementing a robust firewall system

***Using a Cloud Access Security Broker (CASB)**

Regularly updating antivirus software

Conducting frequent penetration testing

Explanation:

Using a Cloud Access Security Broker (CASB) is the correct answer. A CASB acts as a gatekeeper, extending an organization's security policies into the cloud storage infrastructure. It provides visibility of the company data, ensures regulation compliance, controls user access to prevent threats, and secures data through encryption and loss prevention.

While a firewall is an essential part of any security infrastructure, it primarily protects the network from unauthorized external access. It does not extend an organization's security policies into the cloud or provide visibility of the company data in the cloud.

While regularly updating antivirus software is a good practice for protecting against malware, it does not specifically address the security needs of cloud storage.

Penetration testing is a valuable tool for identifying vulnerabilities in a system. However, it does not provide the continuous monitoring and policy enforcement that a CASB does.

q_cloud_stor_cloud_secp8

Cloud storage is a virtual service, so the infrastructure is the responsibility of the storage provider. Access control should be set as a local file system would be, with no need for the provider to have access to the stored data.

You are implementing the following measures to secure your cloud storage:

Verify that security controls are the same as in a physical data center.

Use data classification policies.

Assign information into categories that determine storage, handling, and access requirements.

Assign information classification based on information sensitivity and criticality.

Which of the following is another security measure you can implement?

Answers:

***Dispose of data when it is no longer needed by using specialized tools.**

Create versioned copies of your cloud data.

Configure redundancy and distribution of data.

Configure distributed resources to act as one in a federated architecture.

Explanation:

Disposing of data when it is no longer needed by using specialized tools is another security measure you can implement.

Creating versioned copies of your cloud data, configuring redundancy and distribution of data, and configuring distributed resources to act as one in a federated architecture are all measures that improve the fault tolerance and durability of your data.

q_cloud_stor_risk_secp8

A tech company plans to launch a new application on a cloud platform to cater to its growing customer base. The lead security analyst examines potential vulnerabilities to ensure the application remains secure after deployment.

The analyst focuses on potential weak points within the application's design and the cloud platform's infrastructure.

Considering vulnerabilities associated with applications and cloud platforms, which of the following issues poses the highest risk related to unauthorized data access in cloud-hosted applications?

Answers:

Inefficiently allocated cloud resources

Disabled logging features within the application

***Misconfigured cloud storage access controls**

Absence of multi-factor authentication (MFA) for application users

Explanation:

Misconfigured access controls on cloud storage can allow unauthorized users to access, modify, or delete data. Misconfigured access controls pose a direct threat to data integrity and confidentiality.

While inefficient resource allocation can lead to performance issues and increased costs, it does not directly correlate with unauthorized data access.

While logs are crucial for monitoring and detecting malicious activities, their absence primarily hampers post-incident investigations rather than directly causing unauthorized data access.

While MFA provides an added layer of security, its absence primarily affects user account security. However, a misconfigured cloud storage can allow unauthorized data access even with MFA.

10.5 Mobile Devices

As you study this section, answer the following questions:

Which process allows you to define specific apps that users can have on mobile devices?

Which two configurations can be used to deploy Windows Intune?

What does a mobile device management (MDM) solution allow you to do?

How do jailbreaking and sideloading differ?

In this section, you will learn to:

Enforce security policies on mobile devices.

Sideload an application.

The key terms for this section include:

Term	Definition
App whitelisting	The process of identifying apps that users are allowed to have on mobile devices.
Geotagging	The process of embedding GPS coordinates within mobile device files, such as image or video files created with the device's camera.
Data exfiltration	The unauthorized copy, transfer, or retrieval of data from a computer, server, or network.
Sandboxing	The isolation of an app so that it can't affect other areas of a computer or network.
Jailbreaking	The process of removing inherent protections placed by the device manufacturer.
Sideloading	Installing an app on a mobile device via a method other than the manufacturer's app repository.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
TestOut Security Pro	2.2 Harden network devices 2.2.6 Bring Your Own Device (BYOD) security
CompTIA Security+ SY0-701	2.3 Explain various types of vulnerabilities. Mobile device Side loading Jailbreaking 2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Application allow list Hardening techniques Encryption

3.3 Compare and contrast concepts and strategies to protect data

- Methods to secure data
  - Encryption

4.1 Given a scenario, apply common security techniques to computing resources.

- Secure baselines
  - Establish
  - Deploy
- Hardening targets
  - Mobile devices
- Wireless devices
  - Installation considerations
- Mobile solutions
  - Mobile device management (MDM)
  - Deployment models
    - Bring your own device (BYOD)
    - Corporate-owned, personally enabled (COPE)
    - Choose your own device (CYOD)
  - Connection methods
    - Cellular
    - Wi-Fi
    - Bluetooth
  - Application security
    - Code signing
  - Sandboxing

4.5 Given a scenario, modify enterprise capabilities to enhance security.

	Operating system security Group Policy 4.6 Given a scenario, implement and maintain identity and access management. Password concepts Password managers

10.5.1 Mobile Device Connection Methods (Lesson Video)

Transcript:

Mobile devices pose a unique threat to modern networks. The days when you only had to worry about static workstations that sit in a cubicle are gone. Instead, security professionals also have to worry about countless mobile devices that can be brought in and out of office and are subject to a host of security vulnerabilities.

In this lesson, we're going to take a look at some considerations you need to keep in mind when it comes to mobile devices and the threats they pose to a network.
Let's first take a look at the various connection methods used by mobile devices. Mobile devices use a variety of ways to connect to the internet, the network, and even other devices.
You're probably familiar with cellular and Wi-Fi connections. This is something practically all mobile devices are equipped with. However, mobile devices can also use Bluetooth, NFC, ANT, and infrared to communicate wirelessly.
And each of these connection methods have their own set of security concerns you should be aware of.
For example, Bluetooth is susceptible to Bluejacking and Bluesnarfing, and NFC is susceptible to relay attacks, data corruption, and eavesdropping. So, when implementing mobile devices, be sure to account for the various vulnerabilities each of these connection methods have.
One way to protect against these is to simply disable these connection methods when they aren't being used. This is probably the best way to protect against attacks that leverage vulnerabilities in these systems.
Another way to secure mobile device communications is to employ proper security settings and ensure the use of security protocols. For example, make sure to only allow connections with trusted devices.
Another thing to consider with mobile devices is that of geolocation settings.
Most mobile devices are equipped with GPS, allowing them to use geolocation services. This helps with certain apps, such as maps and weather applications. However, you can also use the GPS unit in mobile devices for security purposes.
For example, using a technique called geofencing, you can lock mobile devices to a particular geographical area. You specify a geographic area where the mobile device is allowed to function a virtual boundary of sorts. If the mobile device moves outside of that area, it can either instruct the user to move back into that area, or even lock the user out from accessing the device altogether.
This functionality can be useful if you issue mobile devices to employees, but don't want them to be able to take them home to use.
Geofencing can also be used to trigger alerts or notifications. For example, it can be used to remind a user to carry out a particular task when they reach a geographic location.
Geolocation services can be of great value to mobile devices, but they also create some security concerns. One of these concerns is that of geolocation data embedded in images.
See, when someone takes a photo on their phone or mobile device, the geographic location where that photo was taken is embedded in the image's EXIF data. If the user was to upload this photo to the internet, a malicious individual could use the EXIF data to carry out a social engineering attack.

To combat this, be sure to disable geotagging on mobile devices.
Another concern with mobile devices is the fact that they are mobile. Because users are able to take them home, to the coffee shop, to wherever, they have the added risk of being lost or stolen. And if that device contains sensitive information or company documents, you risk compromising that information.
To protect against this threat, you can do a couple of things. First, you can employ full device encryption on all mobile devices and require them to have some type of screen lock and authentication method, whether it's a pin, password, biometric, whatever. Just make sure the device requires authentication for access.
Another thing you can do is require all mobile devices to have some type of remote wipe feature. Most mobile devices have built-in remote wipe features, but you might also have to install a third-party remote wipe feature if necessary.
And lastly, mobile devices pose a risk for organizations in the form of data exfiltration.
For example, a disgruntled employee could use their phone to take photos of sensitive information or transfer documents to their phone via Wi-Fi, Bluetooth, or USB. They could even use the device to record sensitive conversations.
If this is a concern for your organization, consider having an extremely strict mobile device policy or only allow the use of company-owned mobile devices on the premises.
Mobile devices pose a unique threat to an organization. And mobile device use is only going to increase in the future. Because of this, it's important that you understand the threats that mobile devices pose to an organization and the steps you can take to properly protect your network.

10.5.2 Mobile Device Connection Facts

Whether it is a tablet, smartphone, or e-reader, all mobile devices share some common characteristics.

This lesson covers the following topics:

- Connection types
- Security considerations
- Application management
- Apps security issues

Connection Types

Mobile devices can use various methods to connect to the internet, network, and other devices. Connection methods include:

- Cellular
- WiFi
- Bluetooth
- NFC
- ANT
- Infrared
- USB
- SATCOM (satellite)

Security Considerations

Mobile device security considerations include:

- Device content management.
- Remote wipe.
- The ability to restrict the device to a particular geographical area (geofencing).
- The ability to manage location information (known as geolocation).
- The requirement to lock the screen with passwords.
- The management of push notification services that can send messages and information to the device when the screen is locked, and the application is not active.
- The ability to store and manage passwords for networks, websites, etc.
- Biometrics.
- Full device encryption.

Application Management

Security considerations regarding the management of applications on the mobile devices include:

- Rooting/jailbreaking/sideloading to load apps from third-party app stores or other websites.
- Flashing with custom firmware.
- Carrier unlocking. This is the ability to use different mobile carrier networks.
- The ability to receive over-the-air (OTA) firmware updates and app updates.
- Camera usage and geolocation information in pictures.
- Text and multimedia message protocols (SMS/MMS).
- Connection to external media.
- Connection using USB OTG (on-the-go).
- The use of a microphone for recording purposes.
- *Tethering* . This is the ability to share internet connectivity with other devices.

Apps Security Issues

When working with mobile device apps, be aware of the following security issues:

Security Issue	Implementation
App control	Application control is implemented in a similar manner for most mobile device operating systems. For iOS devices, all apps come from Apple's App Store, which uses the following mechanisms to secure apps: Running apps are sandboxed. This means they cannot access data stored by other running apps, nor can they access system files and resources. All iOS apps must be digitally signed by Apple or by a third-party developer using an Apple-issued certificate. This ensures that apps from the App Store haven't been tampered with. App developers can use encryption APIs to protect app data. Data can be symmetrically encrypted using AES, RC4, or 3DES. For Windows RT devices, all apps come from Microsoft's Windows Store. The following mechanisms secure apps: Windows RT refuses to load modules not digitally signed by Microsoft. This ensures that apps from the Windows Store have not been tampered with. All apps available through the Windows Store use the Windows RT API, which contains significant security enhancements, including: Windows anti-buffer-overflow memory restrictions Data Execution Prevention (DEP) Address Space Layout Randomization (ASLR) SafeSEH, sacrificial canary values Be aware, however, that iOS devices can be jailbroken. *Jailbreaking* allows apps to be installed from sources other than the App Store. Likewise, apps that aren't from the Windows Store can be installed on Windows RT devices using a process called *sideloading* . Either of these actions can seriously compromise the security of the device and should be avoided. Apps for the Android operating system are not as tightly controlled as those for iOS and Windows RT. Some Android app stores implement good security and tightly controlled apps, much like the App Store and the Windows Store, but others do not. It is strongly recommended that you use apps that come only from a reputable source, such as the following: Google Play Store Amazon Appstore for Android

	Samsung Apps
Authentication and credential management	The average end user must remember passwords for various network resources and services, including web-based services. To make life easier, the credential management function implemented in most mobile operating systems can store usernames and passwords for the end user. A good example is Credential Manager in Windows RT. The iOS operating system performs a similar function using an encrypted keychain for storing digital identities, usernames, and passwords. When the user accesses a password-protected network resource or website, the credential management software supplies the necessary username and password, effectively allowing the user to log in automatically. While credential management software is convenient for the end user, it can also represent a security risk. For example, suppose a user has stored credentials to a sensitive network resource or website on a mobile device and then loses that device. If the user fails to secure the device with a password or PIN, a malicious individual could exploit the stored credentials to gain unauthorized access. It is recommended that you train users not to store credentials to sensitive network resources on their mobile devices.
App whitelisting	App whitelisting is the process of defining specific apps that users can have on their mobile devices. For example, Windows RT provides a feature named Assigned Access, which allows you to define a whitelist of Windows Store applications. Assigned Access ensures that the device has installed only the apps required for its intended purpose. Apps that aren't on the whitelist are not allowed. For iOS and Android devices, app whitelists can be defined and enforced using a mobile device management (MDM) solution.
Geo-tagging	Geo-tagging embeds GPS coordinates within mobile device files, such as image or video files created with the device's camera. While this feature can be useful in some circumstances, it can also create security concerns. For example, if a user regularly posts geo-tagged images to a social media site, cybercriminals could analyze the images and quickly discover where the user works, even the cubicle location. The attacker could possibly derive the user's work hours and personal habits, such as restaurants the user visits for lunch. All of this information could be compiled and used for social engineering attacks on the organization. As a consequence, it is recommended that this functionality be disabled on mobile devices you manage.

10.5.3 Enforcing Mobile Device Security (Lesson Video)

Transcript:

Let's look at mobile device application security.

You need to be aware of several key security issues related to mobile device apps. These issues include making sure that the apps haven't been tampered with, that they don't contain malware, and that they're not sending information to other parties without approval.

The first issue we're going to look at is application control. Let's start by looking at iOS devices.

All apps on iOS devices come from Apple's app store. The Apple app store uses several different mechanisms to secure those apps. One mechanism is sandboxing. A sandboxed app cannot access data stored by other apps running on the device. They cannot access system files or system resources.

In addition, all iOS apps are digitally signed by Apple or a third-party developer. The digital signature uses an Apple-issued certificate and ensures that the apps from the app store haven't been tampered with. App developers can use encryption APIs to protect their app's data. Developers can use symmetric encryption, generated using AES, RC4, or Triple DES.

For Windows mobile devices, all apps come from the Windows store. Like iOS apps, Windows Store apps are digitally signed by Microsoft before they can be loaded on Windows mobile devices. Again, this ensures that the apps from the Windows store haven't been tampered with.

All apps available through the Windows store use the Windows RTAPI, which contains significant security enhancements, such as anti-buffer overflow memory restrictions, data execution prevention (DEP), address space layout randomization, Safe Structured Exception Handling (SafeSEH), and sacrificial canary values.

The Android operating system is based on the Linux kernel and is released through open source licenses.

Some Android stores implement good security and tightly control their apps, much like the App store and the Windows store do. However, other Android app stores do not. If your users have devices with the Android operating system, it is strongly recommended that you require them to use apps that come from a reputable source, such as Google Play Store, Amazon App Store for Android, and Samsung Apps.

Now, even though Apple, Microsoft, and Android control which apps are installed on their devices, it is possible to circumvent these protections by either jailbreaking the device or sideloading an app.

Jailbreaking a device is the process of removing inherent protections placed by the device manufacturer. In the case of an Apple or Android phone, the operating system can be jailbroken in order to allow the installation of non-Apple store apps or to gain elevated root privileges.

In addition, Windows and Android phones are susceptible to app sideloading. With Windows devices, this is when non-Windows store apps are installed on the Windows device. These apps have not been signed by Microsoft and could contain malware. With Android devices, you sideload by installing an application via an APK that was downloaded from the internet.

Authentication and credential management are also key issues with mobile devices.

The average user has to remember usernames and passwords for a variety of network resources and services, including web-based resources. To make life easier, many mobile device operating systems implement some type of credential management system that automatically stores user names and passwords for the user.

When a user accesses a password-protected network resource or website the first time, credential management stores the credentials for that resource. When the user accesses that same resource again, the credential management software automatically supplies the necessary user name and password. This allows the user to automatically log in to that resource. Windows RT has a credential manager. iOS also has a similar function that uses an encrypted key chain for storing digital identities, usernames, and passwords. Android devices can use the Google ecosystem or a third-party app.

Using credential management software is really convenient for the user. From a security standpoint, it presents risks. For example, suppose a user has stored credentials to a sensitive network resource or website on their mobile device, and they lose that device. Any individual who finds the device could exploit that information to gain unauthorized access, especially if the user failed to secure the device with a password or pin number.

Therefore, as a best practice, train your users not to store credentials to sensitive network resources on their mobile devices.

App whitelisting is another mobile device application security consideration. App whitelisting defines specific apps a user can install on a mobile device.

Windows RT provides a feature called Assigned Access. Assigned Access allows you to define a list of allowed Windows Store applications. The mobile device can only have the specified apps installed. For example, the mobile device could be a point of sales system or an educational device. For those purposes, you don't want Angry Birds installed on the device. Any app that's not found on the whitelist is not allowed on the device.

On iOS and Android devices, you can use a Mobile Device Management, MDM, solution to define and enforce an application whitelist. You can use an MDM solution, such as Windows Intune, to provide application whitelisting for Windows RT devices.

Another key issue related to mobile device apps is geolocation data. Geolocation embeds GPS coordinates within mobile device files, such as an image or video file, created with the device's camera.
If geolocation is enabled, the device's GPS location at the time the file is created is embedded within the device files. This includes all pictures taken, such as those on your vacation, with co-workers at lunch, or in your cubicle.
This feature can be useful in some circumstances, but it can also create security concerns. If these pictures or videos are posted on the internet, an attacker can obtain information about you that you don't want them to have.
For example, let's say a user in your organization regularly posts images with geolocation embedded to a social media site like Facebook. A cybercriminal could analyze these images and quickly discover information, such as where the user works, who they work for, and where the user's cubicle is located. A cybercriminal could even derive the user's work hours, work projects, and personal work habits, such as the restaurant they go to for lunch.
All of this information can be compiled and used for social engineering attacks against the organization. Therefore, it is best practice to require that users disable the geolocation setting for mobile devices.
In this lesson, we talked about mobile device application issues. We first talked about making sure you get apps from a reputable source. We talked about credential management and authentication. We looked at application whitelisting. Then we ended this lesson by talking about geotagging.

10.5.4 Enforcing Mobile Device Security Facts

One of the key problems associated with managing mobile devices is the fact that they cannot be joined to a Windows domain. This means Group Policy cannot be used to push security settings to mobile devices automatically.

This lesson covers the following topics:

- Mobile Device Management (MDM)
- Windows Intune
- Windows Intune configurations
- System configuration for Windows Intune

Mobile Device Management (MDM)

One option you can use instead of Group Policy is Mobile Device Management (MDM). Its security settings include the following:

- Security settings can be manually configured on each device. This option doesn't require any additional infrastructure to be implemented. However, it can be a time-consuming task for the administrator (especially in a large organization with many mobile devices) and is not recommended.
- For devices running Apple's iOS operating system, security settings can be distributed in a configuration profile for users to install. The profile can be defined so that only an administrator can delete the profile, or you can lock the profile to the device so that it cannot be removed without completely erasing the device. This option also doesn't require any additional infrastructure for implementation. However, it does rely on the end user to implement the profile, which can be problematic. Additionally, it is not a dynamic strategy, so making even the smallest change to your Mobile Device Security policies would require a great deal of effort to implement.
- A Mobile Device Management solution that pushes security policies directly to each device over a network connection can be implemented. This option enables policies to be remotely enforced and updated without any action by the end user. Many companies have MDM products, including Apple, Cisco, and Microsoft.

Windows Intune

One widely used MDM solution is Windows Intune, which provides cloud-based Mobile Device Management that allows you to remotely manage and secure mobile devices (as well as standard desktop systems starting with Windows 7 or later). Intune cannot be used to manage Windows Server.
The table below shows which operating systems Windows Intune currently supports:

Device	Supported
Apple	Apple iOS 8.0 and later Mac OS X 10.9 and later
Windows	Windows 10 (Home, S, Pro, Education, and Enterprise versions) Windows 10 Mobile Windows Phone 8.1 Windows 8.1 RT PCs running Windows 8.1 Devices running Windows 10 IoT Enterprise (x86, x64) Devices running Windows 10 IoT Mobile Enterprise Windows Holographic & Windows Holographic Enterprise
Google	Google Android 4.0 Google Android for Work

Customers with Enterprise Management + Security (EMS) can also use Azure Active Directory (Azure AD) to register Windows 10 devices.

Windows Intune Configurations

Windows Intune can be deployed in two different configurations:

- Intune Standalone is the recommended deployment method. Intune Standalone is a cloud-only solution that is managed using a web console that can be accessed from anywhere with internet access.
- Hybrid MDM with Configuration Manager is a solution that combines Intune's Mobile Device Management capabilities into Configuration Manager. It uses Intune for policies, profiles, and applications for devices, but it uses Configuration Manager to administer content and manage the devices.

This course covers only cloud-only deployments. Deploying Intune in United Configuration Mode requires experience and skill beyond the scope of this course.

You must first sign up for an account at Microsoft's website before you can use Intune. After you sign up for an Intune account, you can manage the deployment using the following Intune Management Portals:

Account Portal (https://account.manage.microsoft.com) is used to manage subscriptions, users, groups, and domains. End users can also use the account portal to manage their passwords.

Admin Portal (https://admin.manage.microsoft.com) is used to manage enrolled devices and policies.

Company Portal (https://portal.manage.microsoft.com) is used by end users to manage their own accounts and enroll devices.

System Configuration for Windows Intune

After signing up for a Windows Intune subscription, you need to configure the system by completing the tasks listed in the table below:

Configuration Task	Description
Add Intune users	Windows Intune uses administrative and standard users. The first user account created when you sign up for an Intune subscription is made an administrator by default. Additional standard users can be created and managed using the account console by selecting **Management** > **Users** > **New** > **User** . You can also synchronize users and groups into the account console from your Active Directory domain.
Define Intune policies	Intune policies allow you to manage your mobile devices. You can perform tasks such as: Configuring security settings. Applying updates. Configuring firewall settings. Policy settings can be applied to both standalone and domain-joined devices. However, policy conflicts can occur with domain-joined devices. To prevent this from happening, verify that domain-joined devices are not configured to receive the same configuration settings from both Active Directory Group Policies and Windows Intune. Intune provides the following policy templates containing recommended settings that you can deploy: Mobile Device Security Policy Windows Firewall settings Windows Intune Agent settings

	Windows Intune Center settings To set up your Intune policies, access the admin console and select **Policy** > **Add Policy** . Select the policy you wish to deploy and select **Create and Deploy a Policy** . At a minimum, it is recommended that you deploy all of the above policies using the default settings and apply them to either all devices or all users. If necessary, you can later modify the default settings in the policy. You can also configure specific devices or users that a policy applies to.
Manage users and groups	Windows Intune uses two types of groups: User groups allow you to deploy software and Mobile Device Security policies to specific user accounts. Device groups allow you to deploy software, Intune agent settings, and firewall settings to specific devices. To add groups using the Account Console, select Admin > Security Groups > New > **Group** .
Enroll computers	You can enroll standard computer systems (desktops and notebooks) in Windows Intune in one of two ways: Administrator enrollment requires an Intune administrator to set up the enrollment for a specific user. User enrollment allows a user to enroll a computer through the Company Portal. Before you can enroll a system in Intune, you must first download and install the Intune client software on the computer. To do this using administrator enrollment, complete the following: Open a browser and access the Admin Console. Select Administration > Client Software Download > **Download Client Software** . Once this zip file has been downloaded, extract its contents and run the **Windows_Intune_Setup.exe** file as an administrator user. After the installation is complete, restart the computer. The newly managed computer should appear in the Intune Admin Console after a few minutes. Administrator-enrolled computers must be manually linked to an Intune user ID. In the Admin Console, go to Groups > All Devices; then select the device and **Link User** . Users can self-enroll a computer by opening a browser, accessing the company portal, and logging in using their Intune user ID. Then, they select the option to enroll the current device. User-enrolled devices are automatically linked to the user ID that enrolled them.
Enroll mobile devices	To enroll mobile devices in Intune, you must first enable Mobile Device Management in the Admin Console. Select the Administration workspace; then select Mobile Device Management > **Set Mobile Device Management Authority** > **Yes** .

	Users must configure mobile devices with the address of the Intune enrollment server (enterpriseenrollment-s.manage.microsoft.com) during the enrollment process. Be sure users are provided with this address prior to starting the device enrollment process. At this point, Windows RT mobile devices can be enrolled with Windows Intune. To enroll a Windows RT device, search for and run **Company Apps** , then enter your Intune user ID and password along with the address of the enrollment server. Once enrolled, select the link displayed to install the management app from the Windows store. If you want to enroll other types of mobile devices, you must configure Intune for each platform you plan to support. For example, if you want to manage iOS devices, you must obtain an Apple Push Notification service (APNs) certificate and then upload it to Intune. Alternatively, if you plan to support Windows Phone 8 devices, you must get a Windows Phone Dev Center account and upload a signed enterprise mobile code certificate to Intune.

10.5.5 Enforcing Security Policies on Mobile Devices (Demo Video)

Transcript:

There are a lot of things you can do to secure your infrastructure. This is thanks to Microsoft's Intune. In this demonstration, we'll take some steps to secure our environment with Intune. We're going to take a look at Intune security policies, specifically some nice baselines that we're provided with that can really assist us in this manner.

So, what we do is search All Services here and do a search on baselines. This brings us to security baselines. There are these great Microsoft security baselines that we can call on to enact good, secure practices inside of our enterprise. These baselines include some basic security settings that we want to consider.
For example, there's a Microsoft Edge baseline here. We can go ahead and incorporate this into the security posture of our environment. Thanks to Profiles, we can link it in. So let's go ahead and take a look at that. Under Profiles, notice there are no profiles set up for this baseline. But we can easily change that. We'll go in and say this is my Edge profile. We could give it a description if we wanted to. It's automatically limited to Windows 10 and later. As of October 2019, that's Edge version 77 and later. So we'll say Next. Edge has a lot of settings. The defaults are to enable the supported authentication schemes and then list which ones those are. We can set those.
There's the Adobe Flash settings, extension controls, and the ability to save passwords in the software built into Edge. We have many security settings. Notice they're preset a certain way to give a good baseline of security when it comes to the Edge browser. Once we manipulate those, we can then set our scope tags and indicate where we want to assign them. We could go in and set it up for certain groups of users that we've defined, or we could always say we want to make sure this gets applied to every user and every device. The last step is to create this profile. So now we have this nice Edge profile based on this good Microsoft Edge security baseline that we have, and we can deploy this out to our systems.
Notice we've done that by assigning it to all of our devices. After the next sync of our clients, as far as their Edge settings go, they'll be in compliance with the security baseline that we borrowed from Microsoft. That's the one that Microsoft made available to us.
And remember, there were other security baselines for you to look at, too. There's the Windows 10 security baseline, the Defender Advanced Threat Protection baseline, and then of course the Microsoft Edge baseline that we took a look at in this demonstration. Notice that these have versions, too; you're going to have situations where you get new additions of them from Microsoft, and the version increments at that point.

10.5.6 Sideload an App (Demo Video)

Transcript:

In this demonstration, we'll walk through the process of sideloading an app on a Windows 10 workstation system.

Sideloading is the act of installing apps that we create or own on a Windows system without going through the Microsoft Store. This is a significant change from earlier versions of Windows, such as Windows 8 and Windows 8.1.
In Windows 8 and 8.1, sideloading an app was discouraged, and was therefore very difficult to implement. On Windows 10 that has all changed. Sideloading is relatively easy now. With that in mind, let's walk through the process that you need to follow in order to sideload an app.
The first thing we must do is configure Windows to allow sideloading because sideloading is typically not allowed by default. To do this, you right-click Start and select Settings app. Then go to Update and Security > For Developers. Here you have three options.
First, there is Microsoft Store apps, then Sideload apps, and finally Developer mode. This is already set to allow Sideloading apps, so we don't have to change anything. On some versions of Windows, the top option, Microsoft Store, is set as the default. When you change to Sideload apps, you get a message warning of the danger of installing an app by sideloading rather than going through the Microsoft Store.
You can also go into Developer mode. It is similar to the sideloading option, but it also includes advanced development features, which we're not going to work with today. We're just going to use the Sideload apps option. Close this window.
There is another way you can enable sideloading an app and that is through the Group Policy Editor. The great thing about Group Policy is that if you do this on the domain, all systems on the domain that are assigned this Group Policy will be able to sideload an app.
Deciding if you want to do that or not depends on the number of systems involved. To show you how to do that, we'll go into the Local Group Policy and configure this. Remember, you typically would do this on the server, but the settings are the same.
Go down to Search, type â€˜gpedit', and then select Edit Group Policy from the list. The Local Group Policy Editor opens. We'll adjust the size so we can see things better.
Under Computer Configuration, expand Administrative Templates, Windows Components and select App Package Deployment. Next, double-click this setting right here, Allow all trusted apps to install.
When the dialog box opens you will see that by default it's set to Not Configured. We can enable it here, if we want. We're not going to enable it because we already did in the Settings App. This is just another way of doing it, especially if your system is part of a domain.
As always, down here under Help, you can read more about what this setting does if you are not sure. Let's exit out of these windows now.
With sideloading enabled, now we can practice sideloading an app. First, you must have an app to sideload. We've already done that on this system. You could create your own app, but who has time for that?
For this demo, we downloaded one and saved it to the hard disk. Key thing here is that we did not use the Microsoft Store. We determined where this app was hosted on the web and downloaded it. Note that this app has an APPX file extension. This file is the app that we're going to sideload.
Before continuing, we want to check the security certificate for this app. We downloaded the app from the internet, so, you want to make sure the digital signature is OK before you install it. Also keep in mind, if this was an app that we created, we'd have to create a certificate and install that certificate.
Without a certificate, Windows will not allow that app to install. We know this because we tried to do it and guess what? It doesn't work. There are probably work arounds to make it work but that is beyond the scope of this demo.
So, let's check the digital signature of this app. We'll right-click and go to Properties and click the Digital Signatures tab. Next, we'll select the certificate in the Signature list box. Click Details. Under Digital Signature Information, it says, "Hey, this signature is OK." That's a positive thing and we should not have any problems with our app. Click Cancel to close the dialog box.
So, now we have enabled sideloading on the system. We've checked the certificate that was used to sign the app and we've verified that the signature on the app file is valid. The next thing we must do is to sideload the app.
Starting with the Windows 10 Anniversary Edition, you can double click the APPX file and install the app. When we do that, the app installer launches. Here, you can see the name of the app, publisher info, version, and some other info.
If we click Install, the app installs. Let's do things the hard way. We'll install this app using a PowerShell cmdlet. The first thing we must do is run PowerShell. Close the installer here.
We'll right-click the Start and then click Windows PowerShell (Admin). We'll elevate privileges and PowerShell opens.
To make things easier, we put the app into a folder on the C: drive named Install. Let's change directories and go to that folder. To do that type â€˜cd c:/Install' and press Enter. Let's do a â€˜dir' command to see the contents of the folder. The APPX file is there. This makes things easier as you will see in a minute.

Now that we're in the right directory, all we do is type 'Add-AppPackage' followed by the name of the package file that we want to install. We type â€˜./' and instead of typing out this long filename, we'll highlight it, select Edit, right-click, choose Copy, come down here and right-click. When we do, it gets pasted in. Hit Enter and wait a minute while the app is installed on the system.
Something happened, but are we sure? Well, let's see if the app is installed. We'll close these windows, click Start, and scroll down to find the app we just installed. Click it and it opens.
If for some reason we don't want this app any longer, we can uninstall it like we do any Microsoft Store app. You find it in the list, right-click, select Uninstall, select it from the list again, click Uninstall again, and the app is gone.
That's it for this demo. In this demonstration we walked through the process for sideloading an app on Windows 10. First, you must have the app that you want to install. Then you must enable sideloading. You verify that the digital signature on the app is valid, and then use the add appx package PowerShell command to install the app file.

10.5.7 Implement Mobile Device Management

10.5.8 Practice Questions (Section Quiz)

q_mbl_dev_conn_cell_secp8

Your organization recently purchased 18 iPad tablets for use by the organization's management team. These devices have iOS pre-installed on them.

To increase the security of these devices, you want to apply a default set of security-related configuration settings.

What is the BEST approach to take to accomplish this? (Select two. Each option is part of a complete solution.)

Answers:

Require users to install the configuration profile.

Join the tablets to a Windows domain.

Configure security settings in a Group Policy Object.

Configure and distribute security settings in a configuration profile.

***Enroll the devices in a mobile device management (MDM) system.**

***Configure and apply security policy settings in a mobile device management (MDM) system.**

Explanation:

A mobile device management (MDM) solution can push policies directly to each tablet device over a network connection. This option enables policies to be remotely enforced and updated without any action by the end user. The tablet devices must be enrolled in the MDM system before the policy settings can be applied.

One of the key problems associated with managing mobile devices is the fact that they can't be joined to a Windows domain. This means Group Policy can't be used to automatically push security settings to mobile devices.

For devices running Apple's iOS operating system, security settings can be distributed in a configuration profile. The profile can be defined so that only an administrator can delete the profile, or you can lock the profile to the device so that it cannot be removed without completely erasing the device. However, this option relies on the end user to install the profile, which can be problematic.

It's also not a dynamic strategy. Making even the smallest change to your mobile device security policies requires a great deal of effort.

q_mbl_dev_conn_geo_tag_secp8

Recently, a serious security breach occurred in your organization. An attacker was able to log in to the internal network and steal data through a VPN connection using the credentials assigned to a vice president in your organization.

For security reasons, all individuals in upper management in your organization have unlisted home phone numbers and addresses. However, security camera footage from the vice president's home recorded someone rummaging through her garbage cans prior to the attack. The vice president admitted to writing her VPN login credentials on a sticky note that she subsequently threw away in her household trash. You suspect the attacker found the sticky note in the trash and used the credentials to log in to the network.

You've reviewed the vice president's social media pages. You found pictures of her home posted, but you didn't notice anything in the photos that would give away her home address. She assured you that her smartphone was never misplaced prior to the attack.

Which security weakness is the MOST likely cause of the security breach?

Answers:

***Geotagging was enabled on her smartphone.**

Sideloaded apps were installed on her smartphone.

Weak passwords were used on her smartphone.

A Christmas tree attack was executed on her smartphone.

Explanation:

Geotagging embeds GPS coordinates within mobile device files (such as image or video files) created with the device's camera. While this feature can be useful in some circumstances, it can also create security concerns. In this scenario, the vice president probably posted geotagged images to her social media accounts. The attacker likely analyzed the images to discover where she lived and then conducted a dumpster dive attack that yielded the sticky note with the vice president's VPN credentials. The best way to remedy this weakness is to simply disable this functionality in the mobile devices you manage.

Sideloaded apps can only be installed if the device administrator has specifically configured the device to allow them, so this is an unlikely cause.

A weak smartphone password is a concern, but this would not be the cause of the exploit if the device were always in the vice president's possession.

A Christmas tree attack is used to fingerprint network devices, not to gather personally identifiable information.

q_mbl_dev_conn_jailbreak_secp8

A software technician presents a forum on sideloading and jailbreaking to a group of new mobile users.

Which of the following points will the technician include in their discussion of the use of jailbreaking? (Select two.)

Answers:

***It is a method used to gain elevated privileges and access to system files on mobile devices.**

***It allows users to install unauthorized applications and customize device appearance and behavior.**

It refers to the installation of applications from sources other than the official application store of the platform.

It does not undergo the same scrutiny and vetting process as those on official application stores.

It provides the ability to share internet connectivity to other devices.

Explanation:

Jailbreaking is a method used to gain elevated privileges and access to system files on mobile devices.

Jailbreaking allows users to install unauthorized apps, customize the device's appearance and behavior, access system files, and bypass restrictions implemented by Apple.

Sideloading refers to installing applications from sources other than the official app store of the platform, such as Google's Play Store for Android or Apple's App Store for iOS.

While sideloading allows for greater software flexibility and choice, it poses significant risks as sideloaded apps do not undergo the same scrutiny and vetting process as those on official app stores.

Tethering (not jailbreaking) is the ability to share internet connectivity to other devices.

q_mbl_dev_conn_lock_secp8

Which of the following mobile device security considerations disables the ability to use the device after a short period of inactivity?

Answers:

Remote wipe

***Screen lock**

GPS

TPM

Explanation:

A lockout (or screen lock) disables the ability to use the device after a short period of inactivity. The correct password or personal identification number (PIN) unlocks the device.

Remote wipe, also known as sanitization, remotely clears specific, sensitive data on a mobile device. This task is also useful if you are assigning the device to another user or after multiple incorrect password or PIN entries. Data encryption also ensures data confidentiality on the device. Voice encryption (on mobile phones) ensures data confidentiality during transit.

Global Positioning System (GPS) tracking can assist in a device's recovery by displaying its current location.

Trusted Platform Module (TPM) is a hardware chip on the motherboard that can generate and store cryptographic keys to check the integrity of startup files and components.

q_mbl_dev_conn_security_considerations_secp8

Which of the following is NOT a common security consideration for mobile devices?

Answers:

Remote wipe

Geofencing

Full device encryption

***Water resistance**

Explanation:

The following is NOT a common security consideration for mobile devices:

While water resistance can be an important feature for mobile devices, it is not a security consideration. It is more related to the physical durability and resilience of the device, not the protection of data or prevention of unauthorized access.

The following are common security considerations for mobile devices:

Remote wipe is a common security feature for mobile devices. It allows an administrator or device owner to remotely erase all data on the device, which can be especially useful if the device is lost or stolen.

Geofencing is a security feature that restricts the use of a device to a specific geographical area. If the device leaves the designated area, certain functionalities can be disabled or alerts can be sent.

Full device encryption is a security feature that encrypts all data on the device. This means that even if someone gains physical access to the device, they cannot access the data without the encryption key (usually a password or PIN).

q_mbl_dev_conn_sideload_01_secp8

A cyber team presents a discussion on the use of sideloading and jailbreaking to a group of board members.

Which of the following BEST describes sideloading? (Select two.)

Answers:

***It refers to the installation of applications from sources other than the official application store of the platform.**

***It does not undergo the same scrutiny and vetting process as those on official application stores.**

It is a method used to gain elevated privileges and access to system files on mobile devices.

It allows users to install unauthorized applications and customize device appearance and behavior.

It provides the ability to share internet connectivity to other devices.

Explanation:

Sideloading refers to installing applications from sources other than the official app store of the platform, such as Google's Play Store for Android or Apple's App Store for iOS.

While sideloading allows for greater software flexibility and choice, it poses significant risks as sideloaded apps do not undergo the same scrutiny and vetting process as those on official app stores.

Jailbreaking is a method used to gain elevated privileges and access to system files on mobile devices.

Jailbreaking allows users to install unauthorized apps, customize the device's appearance and behavior, access system files, and bypass restrictions implemented by Apple.

Tethering (not sideloading) provides the ability to share internet connectivity to other devices.

q_mbl_dev_conn_sideload_02_secp8

An IT administrator observes that a significant number of mobile devices within the organization have applications installed from outside official app stores.

Concerned about the security implications, the administrator decides to assess the vulnerabilities introduced by this practice.

Which of the following BEST describes the process that allows users to install applications on their devices from sources other than official app stores, potentially exposing the device to malware or unauthorized data access?

Answers:

Rooting

***Sideload**

Jailbreaking

Keylogging

Explanation:

Sideloading refers to the process of installing applications from sources other than official app stores. It poses a risk as the sideloaded apps might not undergo the same security checks as those in official stores.

While rooting provides users with elevated privileges on their devices, allowing them to modify the operating system and bypass certain restrictions, it does not specifically pertain to the installation of apps from outside official app stores.

Jailbreaking allows users to bypass restrictions imposed by Apple and install unauthorized apps, but the term "sideloading" more universally describes the act of installing apps from unofficial sources.

Keylogging is a method used by attackers to record a user's keystrokes, capturing everything typed.

q_mbl_dev_conn_types_secp8

Which of the following is NOT a connection type used by mobile devices?

Answers:

WiFi

Bluetooth

NFC

***HDMI**

Explanation:

The following is NOT a connection type used by mobile devices:

HDMI (High-Definition Multimedia Interface) is a type of connection used primarily for transmitting high-quality video and audio from a device to a display, such as a TV or monitor. While some mobile devices may support HDMI output through a special adapter, it is not a common connection type used by mobile devices in the same way as WiFi, Bluetooth, or NFC.

The following are common connection types used by mobile devices:

WiFi is a common connection type used by mobile devices to connect to the internet or local networks wirelessly.

Bluetooth is another common connection type used by mobile devices for short-range connections to other devices such as headphones, speakers, or other mobile devices.

NFC (Near Field Communication) is a connection type used by mobile devices for very short-range connections, often used for things like contactless payments or transferring data between devices by touching them together or bringing them very close to each other.

q_mbl_dev_conn_unauthorized_apps_secp8

An IT security specialist at a mid-size corporation observes a trend of unauthorized apps appearing on company-provided mobile devices. The specialist suspects the employees are either sideloading apps or have jailbroken their devices.

What steps should the security specialist take to verify the cause of the unauthorized applications and to re-establish proper security protocols? (Select two.)

Answers:

***Conduct device audits to identify unauthorized applications and to detect any signs of jailbreaking or sideloading.**

***Implement mobile device management (MDM) policies to restrict unauthorized application installation.**

Purchase new mobile devices to replace all current ones.

Hire external IT consultants to manage mobile device usage.

Implement a credential management policy.

Explanation:

Audits allow the IT security specialist to directly identify whether unauthorized applications are present, which could indicate sideloading or jailbreaking activities.

Implementing MDM policies ensures that only authorized apps, vetted for security and functionality, are on company devices, reducing the risk of security breaches.

The devices themselves are not the problem; the problem is the behavior leading to sideloading or jailbreaking, which could easily recur on new devices without proper management and policies.

Organizations employ external consultants for broader strategic initiatives when internal resources face an overwhelming amount of tasks, not for routine security management tasks.

Credential management is a feature implemented in most mobile operating systems for storing usernames and passwords for the end user. A good example is Credential Manager in Windows RT. However, implementing a policy for credential management will not help with verifying the cause of the unauthorized applications and re-establishing proper security protocols.

q_mbl_dev_conn_white_secp8

Your organization recently purchased 20 Android tablets for use by the organization's management team.

To increase the security of these devices, you want to ensure that only specific apps can be installed.

Which of the following would you implement?

Answers:

App blacklisting

Credential manager

***App whitelisting**

Application control

Explanation:

App whitelisting is the process of defining specific apps that users can have on their mobile devices. Apps not on the whitelist are not allowed to be installed.

Blacklisting apps is the process of defining specific apps that users cannot have on their mobile devices.

The credential manager function that is implemented in most mobile operating systems can store usernames and passwords for the end user.

Application control is implemented by each mobile operating system. It determines how apps are installed and where they come from.

q_mbl_dev_conn_wipe_secp8

A smartphone was lost at the airport. There is no way to recover the device.

Which of the following ensures data confidentiality on the device?

Answers:

***Remote wipe**

Screen lock

GPS

TPM

Explanation:

Remote wipe, also known as sanitization, remotely clears specific, sensitive data on a mobile device. This ensures that whoever has the device cannot see the sensitive data. This task is also useful if you are assigning the device to another user or after multiple incorrect entries of the password or PIN. Data encryption also ensures data confidentiality on the device. Voice encryption (on mobile phones) ensures data confidentiality during transit.

Global Positioning System (GPS) tracking can assist in the recovery of the device by displaying its current location.

A lockout (or screen lock) disables the device's interface after a short period of inactivity. The correct password or personal identification number (PIN) unlocks the device.

Trusted Platform Module (TPM) is a hardware chip on the motherboard that can generate and store cryptographic keys to check the integrity of startup files and components.

q_mbl_dec_sec_config_secp8

Which of the following is the recommend Intune configuration?

Answers:

Hybrid MDM

***Intune Standalone**

Account portal

Company portal

Explanation:

Intune Standalone is the recommended deployment method. Intune Standalone is a cloud-only solution that is managed using a web console that can be accessed from anywhere with internet access.

Hybrid MDM with Configuration Manager is a solution that combines Intune's mobile device management capabilities into Configuration Manager.

The Account portal is used to manage subscriptions, users, groups, and domains.

The Company portal is used by end users to manage their own account and enroll devices.

q_mbl_dec_sec_enroll_secp8

You are the IT manager of a medium-sized company that has recently adopted Windows Intune for mobile device management. The company has a mix of Windows, iOS, and Android devices. You have already added Intune users and defined Intune policies.

What should be your next step in the system configuration process?

Answers:

Create additional administrative accounts

***Enroll mobile devices**

Synchronize users and groups from your Active Directory domain

Deploy software to user accounts

Explanation:

Enroll mobile devices is the correct answer. After adding users and defining policies, the next step in the system configuration process for Windows Intune would be to enroll the mobile devices that will be managed by the system. This allows the defined policies to be applied to these devices.

While it's important to have administrative accounts for managing the system, the question scenario indicates that users have already been added to Intune. The next logical step would be to enroll the devices that these users will be using.

While synchronizing users and groups from your Active Directory domain is a possible step in the configuration process, the scenario indicates that users have already been added to Intune. Therefore, the next step would be to enroll the devices that these users will be using.

While deploying software is an important part of managing devices with Intune, it would be premature to do this before the devices have been enrolled. The next logical step would be to enroll the devices that will be managed by the system.

q_mbl_dec_sec_mdm_01_secp8

Which of the following is a solution that pushes security policies directly to mobile devices over a network connection?

Answers:

Credential Manager

Group Policy

***Mobile device management (MDM)**

Application Control

Explanation:

Mobile device management (MDM) is a solution that pushes security policies directly to each device over a network connection. MDM solutions enable policies to be remotely enforced and updated without any action by the end user. Many companies have MDM products, including Apple, Cisco, and Microsoft.

The Credential Manager function that is implemented in most mobile operating systems can store usernames and passwords for the end user.

Group Policy cannot be used to automatically push security settings to mobile devices. This is because the devices cannot be joined to a Windows domain.

Application Control is implemented by each mobile operating system. It determines how apps are installed and where they come from (App Store, etc.).

q_mbl_dec_sec_mdm_02_secp8

A large financial institution recently adopted a bring your own device (BYOD) policy. It understands the cost and flexibility advantages of this approach but is concerned about the potential security implications.

Specifically, the institution wants to ensure that its sensitive data remains protected even when accessed from or stored on employees' personal devices.

What would be the MOST effective strategy to safeguard data in this context?

Answers:

Regularly update the company's firewall and antivirus software

***Deploy a mobile device management (MDM) solution**

Implement mandatory password changes every 30 days

Conduct regular security training for employees

Explanation:

An MDM solution allows a company to manage, secure, and enforce policies on employees' mobile devices, even if they are personal devices.

While updating the firewall and antivirus software is a good general practice for any organization, it would not specifically address the security risks associated with a BYOD policy.

While regular password changes can enhance security by limiting the effectiveness of stolen or guessed passwords, they do not directly address the risks associated with a BYOD policy.

While security training is an important aspect of any organization's cybersecurity strategy, it would not directly mitigate the risks associated with a BYOD policy.

q_mbl_dec_sec_mdm_03_secp8

A large firm requires better control over mobile users' access to business applications in the cloud. This access will necessitate single sign-on and support for different device types.

Given the need for risk transference and ease of recovery, what solution should the company consider using?

Answers:

***Third-party mobile device management (MDM)**

Custom in-house software

Virtual private network (VPN) connections

Cloud-hosted desktops

Explanation:

Third-party MDM solutions allow single sign-on, device diversity support, and cloud application access control. They transfer some risk to the provider via service level agreements (SLAs) and often feature rapid recovery options.

While it could meet the exact needs of the firm, developing and maintaining in-house software presents risks and complexities. Additionally, it does not facilitate risk transference as the company retains all responsibility.

Virtual private networks (VPNs) secure network connections but do not manage device-level access controls or single sign-on. Moreover, VPNs do not inherently offer risk transference or ease of recovery.

While cloud-hosted desktops can provide a consistent environment accessible from any device, they do not offer the mobile-specific control and risk transference that a third-party MDM solution does.

q_mbl_dec_sec_mdm_04_secp8

An organization implemented a bring your own device (BYOD) policy for employees to use their mobile devices for work-related tasks. The organization's IT department identified concerns about the security risks associated with BYOD.

They determined that employees' mobile devices must meet the security requirements to protect sensitive company data.

Considering the scenario, which of the following measures is the MOST effective way to enhance the security of employees' mobile devices under the BYOD policy?

Answers:

***Using mobile device management (MDM) solutions to centrally control employees' mobile devices.**

Providing employees with company-owned mobile devices.

Restricting all access to company resources from mobile devices.

Enforcing complex passwords for all employee mobile devices.

Explanation:

Mobile device management (MDM) solutions enable the organization to control employees' mobile devices centrally, enforce security policies, perform remote data wipes in case of loss or theft, and ensure devices remain updated and compliant with security standards.

Although company-owned devices can offer better control, they may not always be feasible under a bring your own device (BYOD) policy and can incur significant costs for the organization.

Completely restricting access to company resources may hamper productivity and negate the advantages of the BYOD policy.

While strong passwords are crucial, relying solely on them does not encompass all security concerns related to BYOD. An MDM solution presents a more comprehensive approach.

q_mbl_dec_sec_mdm_cloud_secp8

The IT team of a medium-sized company plans to implement a mobile device management (MDM) solution to enhance security and streamline the management of its growing number of mobile devices.

The company has employees who use various devices, such as smartphones and tablets, for work tasks in and out of the office.

The IT team needs to choose the MOST appropriate deployment model for their MDM solution to ensure seamless device management and data protection.

Which deployment model for MDM provides the highest level of control and security for the company's diverse mobile devices?

Answers:

***Cloud-based deployment**

On-premises deployment with limited network access

Hybrid deployment with minimal cloud integration

Bring your own device (BYOD) deployment with partial control

Explanation:

A cloud-based MDM deployment allows the IT team to manage mobile devices centrally from the cloud, providing a high level of control, flexibility, and security.

An on-premises MDM deployment requires local infrastructure and may have limited access to manage devices outside the corporate network. This approach might not be suitable for employees using mobile devices off-site.

A hybrid MDM deployment combines on-premises and cloud components. While it offers some control, the minimal cloud integration may limit the ability to manage devices effectively, especially when employees are remote.

BYOD deployment allows employees to use personal devices for work, but it may not provide the desired level of control and security needed to protect sensitive data on diverse devices.

q_mbl_dec_sec_policies_secp8

The IT manager has tasked you with configuring Intune. You have enrolled the devices and now need to set up the Intune policies.

Where would you go to set up the Intune policies?

Answers:

***In the Admin portal, select Policy > Add Policy.**

In the Company portal, select Policy > Add Policy.

In the Admin portal, select Management > Policy > Add Policy.

In the Company portal, select Management > Policy > Add Policy.

Explanation:

To set up Intune policies, access the Admin portal and then select **Policy > Add Policy**.

q_mbl_dec_sec_portal_secp8

Which of the following Intune portals is used by end users to manage their own account and enroll devices?

Answers:

Admin portal

Add Intune Users

Account portal

***Company portal**

Explanation:

The Company portal is used by end users to manage their own account and enroll devices.

The Admin portal is used to manage enrolled devices and policies.

Add Intune Users is a configuration task that is completed in the Account portal.

The Account portal is used to manage subscriptions, users, groups, and domains.

q_mbl_dec_sec_settings_secp8

Your organization recently purchased 20 Android tablets for use by the organization's management team. You are using a Windows domain.

Which of the following should you use to push security settings to the devices?

Answers:

***Intune**

Credential Manager

Group Policy

Application Control

Explanation:

Intune is Microsoft's cloud-based mobile device management (MDM) platform that allows a network administrator to remotely manage and secure mobile devices.

The Credential Manager function that is implemented in most mobile operating systems can store usernames and passwords for the end user.

Group Policy cannot be used to automatically push security settings to mobile devices. This is because the devices cannot be joined to a Windows domain.

Application Control is implemented by each mobile operating system. This determines how apps are installed and where they come from.

q_mbl_dec_sec_windows_server_secp8

Which of the following operating systems does Windows Intune NOT support?

Answers:

Windows 10

Apple iOS 8.0 and later

Google Android 4.0

***Windows Server**

Explanation:

The document clearly states that Windows Intune cannot be used to manage Windows Server.

The following operating systems are supported by Windows Intune:

Windows 10 (Home, S, Pro, Education, and Enterprise versions)

Apple iOS 8.0 and later

Google Android 4.0

10.6 Mobile Device Management

As you study this section, answer the following questions:

What are four methods of mobile device management (MDM)?

What are the benefits of implementing mobile application management (MAM)?

What do Windows Information Protection (WIP) policies provide?

How does Intune help you to secure data?

In this section, you will learn to:

Enroll devices and perform a remote wipe.

The key terms for this section include:

Term	Definition
Windows Information Protection	A technology that helps protect against data leakage on company-owned and personal devices without disrupting the user experience.
Network fencing	Location compliance, known as network fencing, allows you to keep devices outside your corporate network from accessing network resources.
Mobile device management	The administration of mobile devices. MDM software generally allows for tracking devices; pushing apps and updates; managing security settings; and remotely wiping the device.
Mobile application management	The administration of applications on a mobile device. MAM software allows a system administrator to remotely install or remove organizational apps and to disable certain functions within the apps.
Enterprise mobility management (EMM)	A combination of MDM and MAM solutions in one package. EMM allows a system administrator to remotely manage hardware and applications on a mobile device.
Unified endpoint management (UEM)	An all-in-one device management solution. UEM allows a system administrator to manage local and mobile devices, including Internet of Things devices.
Bring your own device (BYOD)	The practice of having employees use their own personal mobile devices for business related tasks.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
TestOut Security Pro	2.2 Harden network devices 2.2.6 Bring Your Own Device (BYOD) security 3.2 Implement application defenses 3.2.1 Implement an application allow list

CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Application allow list 3.1 Compare and contrast security implications of different architecture models. Considerations Ease of deployment 4.1 Given a scenario, apply common security techniques to computing resources. Mobile solutions Mobile device management (MDM) Deployment models Bring your own device (BYOD) 4.2 Explain the security implications of proper hardware, software, and data asset management. Disposal/decommissioning Sanitization Destruction Data retention

10.6.1 Mobile Device Management (Lesson Video)

Transcript:

Today, mobile devices are used in more workplaces than ever before. They create all kinds of new possibilitiesâ€”and all kinds of new problems. So managing mobile devices is a major focus for most organizations.

Most employees bring a personal mobile device to work, usually a cell phone, tablet, or laptop. When these personal electronics are used for business activity, they're called BYOD devices, which stands for bring your own device.
Before mobile devices were common, every device that connected to a network was stationary, such as a desktop or printer, so you could add it to a domain and manage it with Group Policy settings.
But it's not that simple anymore. Organizations need to manage mobile devices and enforce policies to secure their data. If an employee copies sensitive data to a device, the organization needs to be sure that the data will be kept safe outside the organizational network. This is where Mobile Device Management policies come in.
We're going to cover the four main types of mobile device management: Mobile Device Management, Mobile Application Management, Enterprise Mobility Management, and Unified Endpoint Management.
Let's start with Mobile Device Management, or MDM. MDM allows IT administrators to remotely manage a mobile device, even if it's a personally owned device used for work.
MDM lets you track the device, push apps and updates, manage security, and even remotely wipe the device.

This solution works well for company-owned devices, but the release of the iPhone in 2007 caused a huge increase in the number personal devices brought to work and used for work-related purposes, such as email.
Employees didn't want their employers to have too much control over their personal devices. This is when Mobile Application Management solutions were created.
Mobile Application Management, or MAM, solutions focus on managing the applications on a mobile device, not the device itself.
Applications that the organization has developed or any application that the organization provides licenses for falls under MAM policies.
The organization can put limitations on these apps, such as disabling copy and paste functions. IT administration can uninstall the app and the data related to it in case the device is lost or stolen.
As more devices were release and their use spread, MDM and MAM solutions evolved to keep up with changing demands and eventually merged into a solution called Enterprise Mobility Management.
One of the main drawbacks of MDM and MAM solutions is that there are so many different types of devices and operating systems. Each device type requires separate management tools and systems.
Enterprise Mobility Management, or EMM, solutions solve this problem. EMM is a combination of MDM and MAM methods into one package solution. EMM is hardware-agnostic and can remotely manage a mobile device's hardware and applications.
One of the more popular solutions for Enterprise Mobility Management is Microsoft's Intune.
Intune is a part of the Microsoft Enterprise Mobility and Security Suite that's accessed through an organization's Azure portal. Intune allows IT administrators to manage mobile devices and apps, control access to data, and ensure that devices comply with security policies.
Microsoft Intune is included with any Enterprise agreement of at least 500 users, and it supports all mobile device types.
I probably don't have to tell you that integrating device security with Azure's Active Directory make managing mobile devices much, much simpler.
Of course, IT administrators need to worry about managing everything else that's on their network, too: desktop computers, printers, Internet of Things devices, and wearables. This need was met by the creation of another solution called Unified Endpoint Management, or UEM.
UEM tools provide a single point of management for all network devices. Essentially, they are a marriage of traditional device management and Enterprise Mobile Device Management solutions.
All right, that's all for this video. In this lesson, we looked at mobile device management solutions, including Mobile Device Management, Mobile Application Management, Enterprise Mobility Management, and Unified Endpoint Management. We also talked about Microsoft's EMM solution, Intune.

10.6.2 Mobile Device Management Facts

The use of mobile devices in the workplace has increased rapidly over the past few years. The management of these devices has become a big concern for system administrators.

Many organizations allow users to bring their own devices and use them for work-related purposes. This practice, known as bring your own device (BYOD), requires the organization to develop a set of policies to manage these devices, which allow the organization to ensure that mobile devices are secured and can be managed remotely. There are four main types of mobile device management solutions.

This lesson covers the following topics:

- Mobile device management (MDM)
- Mobile application management (MAM)
- Enterprise mobility management (EMM)
- Unified endpoint management (UEM)

Mobile Device Management (MDM)

Mobile device management solutions allow IT administrators to remotely manage a mobile device, even if it's a personally owned device being used for work-related purposes. MDM focuses on managing the device itself but not the applications or software.

Mobile device management provides the ability to:

- Track the device.
- Push apps and updates (also known as provisioning the device).
- Manage security settings, such as lock screens, passwords, etc.
- Remotely wipe the device in case it is lost or stolen.

Mobile Application Management (MAM)

Mobile application management solutions focus on managing the applications on a mobile device but not the device itself. Licensed applications or custom-designed apps fall under MAM policies.

Mobile application management provides the ability to:

- Install and uninstall apps remotely.
- Update apps as needed.
- Limit functionality in an app as needed.

Enterprise Mobility Management (EMM)

Enterprise mobility management combines MDM and MAM solutions in one package. These policies allow a system administrator to remotely manage a mobile device's hardware and applications.

As different brands and manufacturers of mobile devices came on the market, the ability to manage them all became more difficult. Enterprise mobility management solutions address this problem by being able to manage multiple types of devices in a single package.

Microsoft's Intune is one of the most popular EMM solutions. Intune is included with any Windows Enterprise agreement of at least 500 users and supports all types of devices. Intune is integrated into the organization's Azure Active Directory, which simplifies device management even more. Intune allows the system administrator to:

- Manage mobile devices.
- Manage mobile apps.
- Control data access.
- Comply with security policies.

Unified Endpoint Management (UEM)

The need to manage many different devices has become an issue for organizations. Devices such as printers, workstations, servers, and others are managed in Active Directory. However, mobile devices need to be managed separately. A recent solution

to this is unified endpoint management. UEM is the next step in device management. These solutions provide a single point for all types of devices, including:

- Workstations.
- Printers.
- Mobile devices.
- IoT devices.
- Wearable devices.

UEM is the joining together of traditional device management and enterprise mobility management solutions.

10.6.3 Enroll Devices and Perform a Remote Wipe (Demo Video)

Transcript:

There are times when it may be necessary to remove company data from a device that is assigned to one of your users. Now, you could jump in a cab and go to the person's house, knock on the door, and demand to see the device, or you could use some of the cool features that Microsoft has created for us. In this demo, we're going to look at how to restart, retire, or wipe devices on a device remotely.

I'm already logged into Endpoint Manager. I'll come up to Devices and go over to the All Devices blade. Then, in my list of devices, select the Windows11-3 device. There is no particular reason why I picked this device for this demo. In your environment, you would obviously pick the one that you want to work with.
So, now we're seeing some of the basic properties for this device. I can do several things from here. Up here, we have Retire. I'll click on it. If you choose this option, it will remove company data from the device that is managed by Intune. It won't delete the user's personal data. There are some exceptions to what is removed, so make sure you read the documentation and understand what's happening before proceeding.
I don't want to do this, so I'll click No.
Now we'll look at how to wipe a device. I'll click on Wipe. Now a wipe is going to set the device back to a factory reset. You might do this if the person using the device has had an awkward dismissal from a company or the device is lost or stolen. You have two choices. You can wipe the device, but keep the enrollment state and associated account. The other option is to wipe even if the device loses power. It also says that the device might not run Windows 10 or later. This is the more extreme of the two. It's probably the right option if the device is lost or stolen. In fact, it's making me a little nervous. I'm going to click Cancel before I accidentally wipe my demo system!
Next, we have Delete a device. This one deletes the device, and you will no longer see it here in Intune. The device will no longer be able to reach company resources, and the company data may be wiped from the device.
We're going to look at the Restart feature, but first, let's talk about a few others. Here, we have Sync. Use this when you want to sync your device. Fresh Start will remove all preloaded Win32 apps. You can choose whether to retain user data or not. And finally, there are several more things you can do remotely. Now let's go back and look at how to restart a device remotely.
Now I'll click on Restart. I get a message here that basically asks if I'm sure about wanting to restart the device. It goes on to say that the users will not be automatically notified of the restart, and they might lose work. To prevent this, make sure your users are informed before you select this option. I know that there is no one using this device, so I'll go ahead and click on Yes here. You can see a message in the upper right-hand corner of the screen that says Initializing Restart.
Now I'll jump over this this other device and see what happens.
Okay. I'm on the Windows client here, and I've been waiting for a few minutes. It did not reboot instantly. I do have this message that says it's going to sign me out and Windows will shut down in two minutes. So, I'll pause the recording, wait for a few minutes, and see what happens.
Oay., That took about two minutes, and now my device is restarting.

That's it for this demo. In this demonstration, we covered the steps on how to restart, retire, or wipe devices on a device remotely in Intune.

10.6.4 Mobile Application Management (Lesson Video)

Transcript:

All organizations have mobile apps that they like to use. Due to the different types of device platforms as well at data security and ownership these apps can be a challenge to manage.

Mobile Application Management, or MAM, allows the organization to overcome these issues. Intune is Microsoft's MAM solution. MAM in Intune refers to the assortment of management features that lets the system administrator publish, push, configure, secure, monitor, and update mobile apps.
Each app in Intune goes through a life cycle. This cycle is Add, Deploy, Configure, Protect, and Retire. Intune provides a full range of tools to help you manage apps during the life cycle.
Intune allows the organization to develop its own app catalog, create its own self-service portal, and remotely manage apps.
There are a lot of apps available to perform different tasks. But it's important that you only use approved apps that your organization has licenses for and has tested to work with their systems.
An app catalog allows the organization to define which apps a user can and can't use. To put together the catalog, the system administrator often shortened to system admin adds the apps to Intune. Once added, security policies are applied to the apps themselves. But security isn't ensured if you use unapproved apps to access organizational data and resources.
When adding apps to the catalog, the system admin can assign them to specific users or groups. For example, tech support usually doesn't need to use the finance app that the accounting department uses. So that app wouldn't show up in their catalog.
One cool feature that was added to Intune is the ability to blacklist an app. These are apps that absolutely shouldn't be installed on a device that's accessing organizational resources. If the user has one of these apps installed, they'll receive an error message that they're unable to access company resources.
In a large organization, it's just not feasible for the network administrator often shortened to network admin to manually push apps out to all the different groups of devices. But a company can create a self-service portal using Intune that makes the distribution of apps easier for everyone.
Users can access the portal by using a web link. The most efficient way, though, is to have them install the Intune company portal app. When the user installs the app, they log in using their work credentials. The app then allows the user to access their company's portal.
When an organization adds an app to Intune, the system admin decides if it's mandatory to install or not. Intune pushes out mandatory app installs to devices within 24 hours.
Using the portal makes managing apps much easier for the system admins.
So, there's a lot that goes into managing the apps for an organization. That's why Intune is so valuable for its ease and simplicity. With users spread all over the world, knowing that you can remotely manage app installs and usage is a great benefit.
All app types except for the line-of-business apps automatically update as needed. Remember that line-of-business apps are those that someone created internally. So, those updates are pushed out manually through the company portal. Once you load updates to Intune, they should update themselves within 24 hours.
When an employee leaves the organization, the system admin needs to ensure that he or she takes away all access to company resources and removes any company data on the user's devices. Intune allows the admin to remotely remove apps and clear all data from the device without affecting the device itself.
That's it for now. In this lesson, we covered the basics of Mobile Application Management and how you can use Intune to facilitate this. We looked at how you can create a catalog of the apps that users need. Then we covered what a self-service portal is and how you can use it to make the management of apps easier. Finally, we reviewed how you can use Intune to remotely manage apps and even remove apps from user's devices.

10.6.5 Mobile Application Management Facts

Mobile application management (MAM) refers to the assortment of management features that lets a system administrator publish, push, configure, secure, monitor, and update mobile apps. The goal is to ensure users have the applications they need at all times while protecting the organization's data within the apps. This can be very challenging due to the wide variety of device platforms and application types. Intune is Microsoft's MAM solution in the Azure cloud.

This lesson covers the following topics:

- Mobile application management
- Intune application life cycle
- App deployment and update methods

Mobile Application Management

Microsoft app protection policies are rules that make sure the company's data is secure within an application. The user cannot move data or perform any action that is prohibited in a policy. Intune mobile device management (MDM) provides app protection policies that enable MAM to protect devices and data. MAM also provides protection through MAM without enrollment (MAM-WE) in Intune MDM. The following table describes the two possible configurations.

Configuration Option	Description
Intune MDM + MAM	Manage apps using MAM and app protection policies on devices enrolled in Intune MDM. In an MDM + MAM implementation, administrators use the Intune console in the Azure portal.
MAM-WE	Manage apps using MAM and app protection policies but with devices enrolled with third-party enterprise mobility management (EMM) providers. Sensitive data can be managed on any device, including personal devices.

App protection can require a PIN to launch an application.

Intune Application Life Cycle

Each app in Intune goes through a life cycle. Intune provides a full range of tools to help manage apps during each phase. The following table describes these phases.

Phase	Description
Add	Add the apps you would like to manage and assign them in Intune. You can add the following app types: Apps from the Windows Store Apps that are line-of-business apps written in-house Apps on the web Built-in apps

Deploy	Assign the app to users and devices you manage and monitor them on the Azure portal.
Configure	Update deployed apps with new versions using Intune.
Protect	Protect company data in deployed apps with conditional access to email and other corporate resources. Conditional access is based on the criteria you set in the app protection policies that lock down actions the users can perform on devices. Examples of locked-down actions include copying data and preventing app installation on rooted devices.
Retire	Remove apps that have reached end of life or become outdated and are no longer used.

App Deployment and Update Methods

The following table describes the three methods available to work with applications throughout their life cycle.

Method	Description
App catalog	An app catalog allows the organization to define the apps a user can and cannot use. Apps can be assigned to specific users and devices via groups to facilitate management. The catalog is configured to make available to specific users and groups only the apps that they have rights to access. An app can also be blacklisted so no user can use it to access company resources.
Self-service portal	In a large organization, it is not feasible for the network administrator to manually push apps out to all users and groups for all devices. Therefore, a company can create a self-service portal using Intune that makes the distribution of apps easier for everyone.
Remote management	All app types, except for the line-of-business apps, automatically update as needed. Updates can be uploaded into Intune, where they can be pushed out to users and updated within 24 hours. Administrators can push out updates for line-of-business apps through the company portal. When an employee leaves the organization, Intune allows the administrator to remotely remove apps and clear all data from the device without affecting the device itself.

10.6.6 Practice Questions (Section Quiz)

q_mdm_emm_secp8

Which of the following mobile device management (MDM) solutions is hardware-agnostic and supports many different brands of mobile devices?

Answers:

MAM

MDM

***EMM**

UEM

Explanation:

Enterprise mobility management (EMM) is the combination of MDM and MAM solutions in one package. EMM solutions are able to manage multiple brands and types of mobile devices in a single package.

Mobile application management (MAM) solutions focus on managing the applications on a mobile device, but do not manage the device itself.

Mobile device management (MDM) solutions allow IT administrators to remotely manage a mobile device even if it's a personally owned device used for work-related purposes.

Unified endpoint management (UEM) is the next step in device management. These solutions provide a single point for all types of devices.

q_mdm_mam_01_secp8

Mobile application management (MAM) provides the ability to do which of the following?

Answers:

***Remotely install and uninstall apps.**

Control data access.

Comply with security policies.

Manage mobile devices.

Explanation:

Mobile application management (MAM) solutions focus on managing the applications on a mobile device but not the device itself. Licensed applications or custom-designed apps fall under MAM policies.

Mobile application management provides the ability to:

Remotely install and uninstall apps.

Update apps as needed.

Limit functionality in an app as needed.

Microsoft Intune allows the system administrator to:

Manage mobile devices

Manage mobile apps

Control data access

Comply with security policies

q_mdm_mam_02_secp8

The IT department of a medium-sized company is exploring various mobile solutions to improve productivity and enable employees to work efficiently on their mobile devices. They aim to choose a solution ensuring data security and seamless integration with the existing infrastructure.

The team has narrowed the options to three potential mobile solutions: mobile device management (MDM), mobile application management (MAM), and corporate-owned personally enabled (COPE). Each solution offers different features and functionalities, and the IT team is assessing which one BEST meets the company's needs.

Which mobile solution focuses on securing and managing the applications installed on employees' mobile devices rather than the devices themselves?

Answers:

***Mobile application management (MAM)**

Mobile device management (MDM)

Corporate-owned personally enabled (COPE)

A combination of mobile device management (MDM) and mobile application management (MAM)

Explanation:

Mobile application management (MAM) focuses on securing and managing applications on employees' mobile devices. It allows the IT team to control app distribution, updates, and data access without directly managing the entire device.

Mobile device management (MDM) manages and secures the entire mobile device, including settings, configurations, and data.

While corporate-owned personally enabled (COPE) provides a degree of application management, its primary focus is on the complete ownership and management of mobile devices. The scenario seeks a solution focused on securing and managing the applications installed on employees' mobile devices.

While a combination of MDM and MAM is possible, the question specifically asks for a solution focusing on application management, making MAM the more suitable choice.

q_mdm_manage_secp8

What is the minimum number of users needed in a Windows Enterprise agreement for Intune to be included?

Answers:

No minimum

100

***500**

1,000

Explanation:

Intune is included with any Windows Enterprise agreement of at least 500 users and supports all types of devices.

q_mdm_mdm_secp8

Mobile device management (MDM) provides the ability to do which of the following?

Answers:

Update apps as needed.

***Track the device.**

Control data access.

Remotely install apps.

Explanation:

Mobile device management (MDM) solutions allow IT administrators to remotely manage a mobile device even if it's a personally owned device being used for work-related purposes.

Mobile device management provides the ability to:

Track the device.

Push apps and updates (this is also known as provisioning the device).

Manage security settings, such as lock screens, passwords, etc.

Remotely wipe the device in case it is lost or stolen.

Mobile application management provides the ability to remotely install and uninstall apps.

Microsoft Intune allows the system administrator to:

Manage mobile devices

Manage mobile apps

Control data access

Comply with security policies

q_mdm_uem_secp8

Which of the following mobile device management (MDM) solutions allows an organization to manage all devices, including printers, workstations, and even IoT devices?

Answers:

MAM

MDM

EMM

***UEM**

Explanation:

Unified endpoint management (UEM) is the next step in device management. These solutions provide a single point for all types of devices. This includes workstations, printers, mobile devices, IoT devices, and wearable devices.

Mobile application management (MAM) solutions focus on managing the applications on a mobile device, but not on managing the device itself.

Mobile device management (MDM) solutions allows IT administrators to remotely manage a mobile device even if it's a personally owned device being used for work-related purposes.

Enterprise mobility management (EMM) is the combination of MDM and MAM solutions in one package.

q_mam_add_secp8

Which of the following is the first phase of the Microsoft Intune application life cycle?

Answers:

Deploy

Configure

***Add**

Protect

Explanation:

The first phase of the Microsoft Intune application life cycle is to add the apps that are to be managed and assigned in Intune.

Deploy is the second phase.

Configure is the third phase.

Protect is the fourth phase.

q_mam_catalog_secp8

Which of the following app deployment and update methods can be configured to make available to specific users and groups only the apps that they have rights to access?

Answers:

Self-service portal

Remote management

BYOD

***App catalog**

Explanation:

An app catalog allows an organization to define the apps that a user can and cannot use. Apps can be assigned to specific users and devices via groups to facilitate management. The catalog is configured to make available to specific users and groups only the apps that they have rights to access. An app can also be blacklisted so no user can use it to access company resources.

A company can create a self-service portal using Intune that makes the distribution of apps easier for everyone.

With remote management, all app types (except for line-of-business apps) automatically update as needed.

Bring your own device (BYOD) is a policy that allows a user to use their personal device for business purposes.

q_mam_deploy_secp8

In which phase of the Microsoft Intune application life cycle would you assign an app to users and/or devices you manage and monitor them on the Azure portal?

Answers:

***Deploy**

Configure

Add

Protect

Explanation:

During the Deploy phase, apps are assigned to users and devices and then monitored on the Azure portal.

The Configure phase is when apps are updated using Intune.

Add is the first phase. This is when apps are added to Intune to be managed.

Protect is the fourth phase. This is the phase in which company data is protected.

q_mam_remote_01_secp8

Which of the following app deployment and update methods allows updates to be uploaded onto Intune where they can be pushed out to users within 24 hours?

Answers:

Self-service portal

***Remote management**

BYOD

App catalog

Explanation:

With remote management, all app types (except for line-of-business apps) automatically update as needed. Updates can be uploaded onto Intune where they can be pushed out to users within 24 hours.

A company can create a self-service portal using Intune. This makes the distribution of apps easier for everyone.

Bring your own device (BYOD) is a policy that allows a user to use their personal device for business purposes.

An app catalog allows an organization to define the apps that a user can and cannot use.

q_mam_remote_02_secp8

Which of the following app deployment and update methods allows an administrator to remove apps and clear all data from a device without affecting the device itself?

Answers:

Self-service portal

***Remote management**

BYOD

App catalog

Explanation:

With remote management, when an employee leaves an organization, an administrator can remotely remove apps and clear all data from a device without affecting the device itself.

A company can create a self-service portal using Intune that makes the distribution of apps easier for everyone.

Bring your own device (BYOD) is a policy that allows a user to use their personal device for business purposes.

An app catalog allows the organization to define the apps that a user can and cannot use.

q_mam_we_secp8

You are the IT manager of a large multinational corporation. The company has recently decided to implement a bring your own device (BYOD) policy, allowing employees to use their personal devices for work purposes.

The company has a diverse range of device platforms and application types. Your task is to ensure that the company's data is secure within applications on these devices, and that sensitive data can be managed on any device, including personal devices.

Which of the following configurations would be the MOST suitable for this scenario?

Answers:

Intune MDM + MAM

***MAM-WE**

App catalog

Self-service portal

Explanation:

MAM-WE is the correct answer. MAM-WE allows for the management of apps using MAM and app protection policies but with devices enrolled with third-party enterprise mobility management (EMM) providers. This means that sensitive data can be managed on any device, including personal devices, making it the most suitable option for a BYOD policy.

Intune MDM + MAM allows for the management of apps using MAM and app protection policies on devices enrolled in Intune MDM. However, this option is not the most suitable for a BYOD policy as it requires devices to be enrolled in Intune MDM, which may not be feasible or desirable for personal devices.

While an app catalog allows the organization to define the apps that a user can and cannot use, it does not provide the necessary security features required for a BYOD policy.

A self-service portal can make the distribution of apps easier, but it does not provide the necessary security features required for a BYOD policy.

10.7 BYOD Security

As you study this section, answer the following questions:

How would you remediate a tablet or phone infected with malware?

What is an acceptable use policy (AUP)? How does it benefit mobile security?

How does virtual desktop infrastructure (VDI) provide enhanced security and better data protection?

What is the difference between choose your own device (CYOD) and corporate owned, personally enabled (COPE)?

How can you prevent malicious insider attacks?

In this section, you will learn to:

Secure mobile devices.

Secure an iPad.

Create a guest network for BYOD.

The key terms for this section include:

Term	Definition
Bring your own device (BYOD)	A BYOD policy allows employees to use personal devices for work related tasks.
Acceptable use policy (AUP)	An AUP determines the rules for using corporate resources, such as internet access, computers, etc.
Virtual desktop infrastructure (VDI)	VDI is a technology that uses virtual machines and virtual desktops.
Choose your own device (CYOD)	In a CYOD system, the company provides a list of approved devices for an employee to choose from. The ownership and management of devices varies by organization.
Corporate owned, personally enabled (COPE)	In a COPE system, the company provides a list of approved devices for an employee to choose from. The company owns the device; the employee uses and manages the device.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
TestOut Security Pro	2.2 Harden network devices 2.2.6 Bring Your Own Device (BYOD) security 3.2 Implement application defenses 3.2.1 Implement an application allow list
CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Application allow list Isolation 3.1 Compare and contrast security implications of different architecture models. Considerations Ease of deployment 4.1 Given a scenario, apply common security techniques to computing resources. Hardening targets

- Mobile devices
  - Wireless devices
    - Installation considerations
  - Mobile solutions
    - Mobile device management (MDM)
    - Deployment models
      - Bring your own device (BYOD)
      - Corporate-owned, personally enabled (COPE)
      - Choose your own device (CYOD)
    - Connection methods
      - Cellular
      - Wi-Fi
      - Bluetooth

4.2 Explain the security implications of proper hardware, software, and data asset management

- Disposal/decommissioning
  - Sanitization
  - Destruction
  - Data retention

4.5 Given a scenario, modify enterprise capabilities to enhance security.

- Network access control (NAC)

4.6 Given a scenario, implement and maintain identity and access management.

- Access controls

5.1 Summarize elements of effective security governance.

- Policies

	Acceptable use policy (AUP)

10.7.1 BYOD Security Issues (Lesson Video)

Transcript:

Let's spend a few minutes talking about managing mobile devices.

When managing a modern computer network, you need to take into account all of mobile devices that will be connected to that network from both a management perspective and from a security perspective. Most users today have more than one mobile device that they will use to complete their day to day work, including smart phones and tablets. These devices may be purchased by your organization, or they may be personal devices brought from home. This practice is sometimes referred to as bring your own device, or BYOD. While these devices allow employees to be more productive, they also create management issues and open up security holes in your network. Because these devices are so well liked, it really isn't feasible for you to try to prevent their usage. Instead, you need to formulate plans to address them in your overall management strategy.

A key issue to remember is the fact that these devices probably contain some type of sensitive information. The degree of sensitivity depends on your organization, the person's role, and the information downloaded to these devices. It is possible that sensitive information can leak out of your organization through one of these mobile devices. You need to define processes for managing these devices. Let's take a look at a few things you need to consider.

First, you should define a process for provisioning users with mobile devices. Remember, whenever you give an employee one of these devices, it can open up a security risk. You need to carefully control who gets these devices and what they put on them. Begin by defining a procedure for provisioning users with mobile devices. This ensures that you know exactly who has a mobile device and what they will use it for.

Next, you should define acceptable uses for these devices. There are several key aspects related to acceptable use; one is personal use, and the other is after normal hours use. Employment terms play a factor. If an hourly employee is issued a mobile device, and that employee uses the device after normal work hours for company purposes, such as answering emails or responding to requests, then that individual could be eligible for overtime pay. This issue also relates to personal use of the device. If an employee uses a company owned mobile device after hours for anything that's irresponsible, illegal, embarrassing, or even malicious, your organization may have a degree of liability. The liability would stem from the fact that the organization issued the device to the employee and didn't specify acceptable personal use guidelines. Therefore, you should include provisions in your AUP that explicitly defines what constitutes acceptable personal use and which activities are prohibited.

If your organization allows personally owned devices, then you must address the question of who owns company data if it gets copied to a user's device. Does the organization still own it, or does the user own it because it's on their personal device? To deal with this problem, your AUP should define the kind of data that is allowed on mobile devices, especially if the device is personally owned. It should also define the kind of data prohibited on mobile devices. You can implement information classification labels to make it clear to the user into which category a particular piece of data is sorted.

In addition, your AUP needs to address the threat of insider attacks. A malicious internal user could use a mobile device to conduct an insider attack. For example, the user could use the built-in camera to take pictures of sensitive internal information. They could copy sensitive data to the device's storage. They could use the built-in microphone to record conversations. They could use the built-in video function to record proprietary processes, and so on. They could then transmit this stolen information outside your organization without going through your organization's network security mechanism using the device's mobile broadband connection. This would bypass your intrusion protection systems, your firewall, and any other network security mechanisms you may have implemented. Essentially, data could be stolen and you will have no idea that it's been leaked or how it got out.

Because these devices have Wi-Fi or mobile broadband networking built-in to them, it is very difficult to prevent data leakage. If the user can download or capture sensitive information to the device, there is no technological means to prevent data from being transmitted. Accordingly, your acceptable use policy should specify where and when employees can possess mobile devices and how they can use them. For example, you could prohibit mobile devices in the high-security areas of your organization. You could specify in the AUP that any personally owned devices brought on site are subject to random searches.

Once your AUP has been defined, you need to develop a management plan for managing the mobile devices in your network, both company owned and personally owned. For example, one issue you need to be concerned with is

malware propagation. If a user's tablet or phone is infected with malware, the infection can easily spread to your organization's network when they connect their device to your organization's internal network.
One solution is to implement Network Access Control (NAC) so that devices can be remediated before allowing them on your network. The NAC can install the latest antimalware software updates to a device and then run a scan to ensure the device is clean before allowing it on your network.
Alternatively, you could implement a guest wireless network that's isolated from your organization's production network. In this configuration, a BYOD device connects to the guest network to gain Internet access, but is quarantined from the rest of your organization's production network. Using a guest wireless network prevents malware on a BYOD device from infecting your production systems.
You also need to consider policies and procedures for protecting data that is legitimately copied to mobile devices within the guidelines set in your AUP. Organizations spend a lot of time, effort, and money building internal security controls to protect their data. However, if a user copies that information to a mobile device, that data is at risk. For example, if the user loses the device, the data on that device may be compromised. All the firewalls, content filters, and intrusion detection systems that you setup on your internal network are completely ineffective at preventing this. If that device doesn't have the appropriate security settings, anyone who picks up that device will be able view the sensitive data.
To circumvent this risk, consider using a mobile device management (MDM) infrastructure, such as the Windows Intune service. You can require all mobile devices, even personally owned devices, to be enrolled in your MDM infrastructure. You can use this infrastructure to enforce mobile device security policies. For example, using an MDM, you can require enrolled devices to use some form of authentication, such as a PIN or (better yet) a password. You don't want just anyone to be able to pick up a mobile device and begin using it without having to authenticate in some way. Without a password or a PIN code, anyone who finds a lost mobile device could immediately access the information that has been saved on it.
Also consider implementing some type of device lockout with your MDM infrastructure. If a device is left on, but there has been no activity for a period time, the lockout feature should lock the device and require authentication to re-access it. For example, suppose I'm at the airport using my Smartphone. I set my phone down to get my luggage and I forget and leave it in the terminal. After a few minutes of inactivity, the lockout feature automatically locks the device and requires some type of authentication to get back in.
You can also use the MDM infrastructure to enable device encryption. Even if I require authentication and lockouts on enrolled mobile devices, there is still a risk that someone could physically open a lost or stolen device, remove its storage, and try to read its contents from a different system. To prevent this, you can enable device encryption. Without the correct encryption key, the contents of the device's storage are undecipherable. Again, this can help protect data on a device if it gets lost.
You can also use the MDM infrastructure to remotely wipe mobile devices. This feature sends out a command over the Internet to a lost or stolen mobile device that causes the contents of the device to be wiped. The ability to perform a remote wipe is essential if you allow sensitive company information to be copied to mobile devices. If the user loses a device, you need to be able to perform a remote wipe as quickly as possible. You can even use the MDM infrastructure to automatically wipe data from a device after a certain number of failed authentication attempts have occurred. If someone finds or steals a device and then repeatedly tries to guess its pin or password, all the data will be removed from the device.
Finally, you should implement some type of reporting process for your mobile devices. Identifying a lost or stolen device as quickly as possible is probably one of the most important components in your overall mobile device management strategy. If a user loses a device, or if it's stolen, the loss needs to be reported immediately. Remember, users will be either hesitant or embarrassed to let you know that the device that the organization paid for and issued to them is now gone. They may fear some type of retribution. In spite of this, you need to make sure they understand that getting the missing device wiped clean of sensitive data is the organization's primary concern. Make sure users understand that timing is critical. The longer a missing device goes unreported, the more time an attacker has to defeat its security mechanisms and the more likely the sensitive data it contains will be exposed.
So, in this lesson we talked about some things you can do to increase the overall security of the mobile devices used by your organization. We talked about implementing a procedure for provisioning mobile devices, defining an acceptable use statement, and using an MDM infrastructure. Then we looked at some general provisions you can use to protect the data on the device, such as using authentication, implementing a lockout, using remote wipe after failed authentication attempts, encrypting the device, and instructing users to report immediately if a device is lost or stolen.

10.7.2 BYOD Security Facts

In addition to mobile devices owned by your organization, you must also take into account personally owned mobile devices that employees bring to work and use to complete daily work-related tasks. This practice is sometimes referred to as Bring Your Own Device (BYOD). Even though it entails a host of security risks, this is a very common practice in the modern work environment.

This lesson covers the following topics:

- BYOD security issues
- Deployment model alternatives

BYOD Security Issues

Security administrators need to keep the following BYOD security issues in mind:

BYOD Issue	Description	Possible Remedies
Malware propagation	If a user's tablet or phone is infected with malware, the infection can be spread when they connect their device to the organization's network.	Consider implementing a network access control (NAC) solution that remediates devices before allowing them to connect to your network. Alternatively, consider implementing a guest wireless network that is isolated from your organization's production network. User-owned devices can connect to this network to gain internet access but are quarantined from the rest of your organization's production network.
Loss of sensitive data control	If a user copies sensitive data to their device, the organization could potentially lose control of that information. Even the question of who owns the data after it has been copied to the personal device becomes problematic. Consider the following scenarios: The user may not have implemented appropriate security settings on their device, allowing anyone who gains access to the device to view the sensitive data. The user may lose the device, allowing anyone who finds it to access the sensitive data. The device may become infected with malware, potentially exposing sensitive data.	Implement an acceptable use policy that defines which kinds of data are allowed on personally owned devices and which kinds of data are prohibited. Information classification labels can be useful when implementing this policy. Consider requiring personal devices to be enrolled with a mobile device management infrastructure, such as Windows Intune, to enforce mobile device security policies.

Malicious insider attacks	If a user is so inclined, they could use their mobile device to conduct a malicious insider attack. For example, they could: - Use the built-in camera, which nearly all modern mobile devices have, to take pictures of sensitive internal information. - Use the built-in microphone to record conversations. - Use the built-in video function to record proprietary processes and procedures. - Use the device's mobile broadband connection to transfer stolen data to parties outside the organization, bypassing the organization's network security mechanisms.	Implement an acceptable use policy that: - Specifies where and when mobile devices can be possessed within the organization. For example, the possession of mobile devices may be prohibited in high-security areas. - Notifies users that personally owned devices are subject to random searches if brought on-site.
Device management	If a user brings a personally owned device on-site, the organization needs to address clearly who is responsible for managing the device. Responsibility for the following needs to be defined: - Operating system updates - App updates - Anti-malware installation - Anti-malware definition updates	Relying on the end user to implement these updates is unwise. Instead, consider implementing a network access control (NAC) solution that remediates devices before allowing them to connect to your network.
Support	If a user brings a personally owned device on-site, the organization needs to clearly address who will provide support for the device and the apps used on the device. Will the organization's help desk provide support, or must the user depend upon support provided by the device manufacturer?	Implement an acceptable use policy that specifies: - Where users can get support for personally owned mobile devices. - Which apps are allowed for use with organizational data. - Where users can get support for these apps.

Deployment Model Alternatives

To better secure mobile devices used by company employees, consider the following deployment model alternatives to BYOD:

Corporate-owned device: A corporate-owned device strategy lets businesses more effectively monitor and control activities performed on mobile devices. One advantage of this model is that businesses can purchase devices at significant discounts. The corporate-owned model also includes the option of restricting mobile device use to the workplace only. However, employees who need access to corporate email and other data after hours may feel compelled to use their personal devices for such access.

Corporate-owned, personally enabled (COPE): The COPE model gives businesses significant control over device security while allowing employees to use the devices to access both corporate and personal data. Because the company owns the device, it can be secured more easily and wiped clean if lost or stolen. One disadvantage of this model is that employees who are not free to choose their own devices may bring their own anyway.

Choose your own device (CYOD): The CYOD model provides slightly more flexibility in giving users a limited selection of devices to choose from. However, since the devices are still corporate-owned, IT managers can implement more effective security measures to prevent breaches.

Virtual desktop infrastructure (VDI): VDI can be used with any of the above models, including BYOD, to allow mobile devices to establish a remote connection to a virtualized desktop. Using VDI provides enhanced security and better data protection because most of the data processing is provided by servers in the data center rather than on the local device.

10.7.3 Securing Mobile Devices (Demo Video)

Transcript:

In this demonstration, we're going to discuss mobile device security. Specifically, we're going to look at a few things you can do to increase the overall security of an iPad running the iOS operating system.

Let's begin with passcode locks. I'm going to go to Settings, then to Touch ID & Passcodes.
Now, I need to enter my current passcode.
You can use passcodes on your iPad to protect the information stored on the device. If you set a passcode, then every time you either power on the device or wake it up, you will be prompted to enter the passcode before you can run any apps or access any of the information on the device.
That's a really good idea in case you forget this device in a taxicab, hotel, train station, or airport. You don't want just anyone to pick up the device and have full access to whatever's on it.
So, let's look at how you do this. Notice that, right now, passcodes are currently turned on. Let's go ahead and tap Turn Passcode Off, and then we're prompted to confirm that. I'll tap Cancel.
Right below that, I can change the passcode. And when I tap it, the iPad wants to know what my current passcode is before it allows me to change to a new one.
So, that is how you work with passcodes.
In addition, we can also specify how long this device can be idle before we're going to require the passcode to be re-entered. So, first, we need to go to Display & Brightness. And then to Auto-Lock right here. Notice that currently, Auto-Lock is set to Never. That means that this iPad can sit idle forever and the user will not be required to re-enter the passcode.
Let's change that. I'm going to tap on Never. And let's set the auto-lock for 15 minutes. Now, if this iPad sits idle for 15 minutes, meaning I'm not tapping anything, then the screen is automatically going to lock.
So, to this point, we've enabled a passcode, and we've set our screen to automatically lock after 15 minutes. So far, we've increased the security a little bit on this device.
But there are more things we can do. I'll tap Touch ID & Password again. I'll enter in my passcode.

Now scroll to the bottom of this screen. There's an option called Erase Data. If I turn Erase Data on, then all the data on this iPad will be erased if the user tries to enter an incorrect password 10 times in a row. After 10 consecutive failed passcode attempts, all the data on this iPad is going to get erased.
Let's turn that option on. And it warns us, "Hey, if you turn this option on, then everything on this iPad is going to be erased after 10 failed passcode attempts." We're going to go ahead and say, "Nope." I'm borrowing this iPad, so the owner might not like that.
This is a powerful option, and you need to carefully decide whether you want it on. If you have it turned on and someone enters the wrong passcode 10 times in a row, you'll have to restore all the data on this iPad from backup, either from iTunes on a PC using a USB cable or pulling it down from iCloud.
Basically, what we're assuming is that if somebody enters the wrong passcode 10 times in a row, it isn't their iPad. So, at this point, we've really increased the overall security of this device.
Another aspect of mobile device security involves being able to track down a lost device because, well, mobile devices do get lost. They get lost all the time. People leave them on airplanes. They get left in hotels, taxicabs, elevators, and train stations. If this happens, you not only want to protect the data on the device, you also probably want to figure out where that device is.
Many mobile devices provide functionality that you can use to locate the lost device. The iOS operating system running on this iPad provides a function called Find My iPad, and it's tied to the iCloud backup function. If we enable this function and then we lose this iPad, we can sign into iCloud from any web browser on a PC or a laptop system and then use the Find My iPad option to view the device's approximate location on a map.
I tapped iCloud a second ago. Now scroll down a little bit. Notice down here that the Find My iPad option is already turned on. This gives us a lot of options.
For example, you could have a message displayed on the iPad that tells whoever finds it how to get in touch with you. You could even use iCloud.com to go in and remotely change the passcode lock assigned to this device. And if you're really concerned about not losing the data on a lost device, then you can use iCloud.com to perform a remote wipe. Essentially, what we're doing with a remote wipe is assuming that the iPad is lost, that the person who found it has no intention of returning it, and that they may want to get at the information that's stored on it. In this case, you can use iCloud.com to send a remote wipe command to the iPad, which will then restore the iPad to its original factory settings. Any personal information on it, as well as any proprietary sensitive information belonging to your organization, gets wiped out.
If you have this option turned on, it's not a bad idea to come over here, under Backup, and turn on iCloud Backup as well. The idea here is that if you ever do end up performing a remote wipe, or if you end up losing the iPad altogether and never get it back, then you can restore all the data from your old iPad onto a new iPad. Go ahead and tap Cancel for now.
That's it for this demonstration. In this demo, we talked about some things you can do to increase the overall security of a mobile device, in this case, an iPad. We began this demonstration by talking about how to set passcode locks. We talked about how to set the auto-lock feature of the iPad. We talked about wiping the data off the iPad if somebody enters the wrong passcode more than 10 times in a row. Then we ended this demonstration by talking about how to locate a lost device using the Find My iPad option and iCloud.

10.7.4 Secure an iPad (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. The receptionist uses an iPad to manage employees' schedules and messages. You need to help her secure the iPad because it contains all of the employees' personal information.

In this lab, your task is to:

View the current iOS version and then answer the applicable question.

Apply the latest software update and then answer the applicable question.

Configure Auto-Lock with a five-minute delay.

Configure Passcode Lock using a passcode of **C@sp3r**

Require the passcode after five minutes.

Configure Data Erase to wipe all data after 10 failed passcode attempts.

Require unknown networks to be added manually.

Turn off Bluetooth.

Explanation

Complete this lab as follows:

Verify the current version of iOS installed on your iPad.

Select **Settings** .

From the Settings pane, select **General** .

From the General pane, select **About** .

From the top right, select **Answer Questions** .

Answer Question 1. Minimize the question dialog.

Apply the latest software update.

From the About pane's heading, select **General** . This returns you to the General settings.

From the General pane, select **Software Update** .

Select **Download and Install** .

Select **Agree** .

Select **OK** . The software is downloaded.

Select **Install** .

The installation automatically starts after 10 seconds.

Slide the **arrow** to the right to unlock the iPad.

From the top right, select **Answer Questions** .

Answer Question 2 and then minimize the question dialog.

Configure Auto-Lock.

From the Settings pane, select **Display & Brightness** .

From the right pane, select **Auto-Lock** and then select **5 minutes** .

Configure Complex Passcode Lock and Data Erase.

From the left menu, select **Touch ID & Passcode** .

From the right pane, select **Turn Passcode On** .

Enter the new passcode of **C@sp3r**

Select **Next** .

Re-enter **C@sp3r** .

Select **Done** .

Scroll down and then slide ***Erase Data*** to **ON** .

Select **Enable** .

Select **Require Passcode** .

Select **After 5 minutes** .

Require unknown networks to be manually added.

From the left menu, select **Wi-Fi** .

Select ***Ask to Join Networks*** , then select **OFF** .

Turn off Bluetooth as follows:

From the left pane, select **Bluetooth** .

Slide ***Bluetooth*** to **OFF** .

From the top right, select **Answer Questions** .

Select **Score Lab** .

10.7.5 Creating a Guest Network for BYOD (Demo Video)

Transcript:

In this demonstration, we're going to lookÂ at creating a bring your own device (BYOD)Â guest wireless network for employees.Â This can be a very useful if you're anÂ enterprise wireless network administrator.

When employees work for an organization, they expect to beÂ able to have internet access for their personal wirelessÂ devices so they can check social media, email, and so on.
We don't want them to use our standard wirelessÂ network because that would be insecure.Â Instead, we can create an isolated wireless network thatÂ has access to the internet but not our internal network.

We're going to use this particularÂ vendor's wireless networking equipment.Â The way you do it with each vendor'sÂ wireless equipment will be different, butÂ the same general principals apply.
Also keep in mind that your organizationÂ should have a policy in place on how theyÂ want to create their wireless networks.
Before we go too far, we could start configuring some of the things on a different device. For example, we could use our network security device to create a VLAN and handle DHCP for us. It all depends on how your organization wants to configure things and have them managed.
Let's continue.
For this vendor's equipment, what we needÂ to do is define a guest access service for our BYOD.Â Then, we need to create a guest wireless LAN.
The first thing we need to do is click Networks.
We want to choose Create New Network.
On this page we decide on a name. I'll call it BYOD, which stands for bring your own device.
Now we're going to choose Guest for our network type. When we pick Guest, that network will be preconfigured with certain rules and policies to help protect our other networks.
Next, we're going to specify that this network will be VLAN 50. This tags the traffic so that our other network equipment knows how to handle it.
Now we need to specify the IP for the BYOD network. I'm going to give it an IP of 192.168.50.1 with a /24 subnet. This can be anything you want, but I like to name things to correlate with the VLAN. So, in this case I made the third octet 50 to match the VLAN. You don't have to, it's just something I do to help me remember that when I see addresses with 50 dot something, it's VLAN 50. This is a Class C private IP range and it'll give us 254 IP addresses that we can use. When I hit Enter, it populated some other settings automatically for me. I have the gateway, the broadcast IP, total IP count, the range, and my subnet mask.
I'll scroll down a bit and then come over here and click on Update DHCP Range. This populates a suggested range of addresses that I can use for my DHCP and hand out to the guest. It put in 50.6 to 50.254, which is fine with me. By the way, I do like to leave a few addresses at the beginning of my networks for potential static assignments. So this worked out well.
I'll scroll down a bit more and now let's look at our DHCP lease time. By default, it's set to 86,400 seconds. For those of you who don't like math, that's 24 hours. Now, I might want to shorten that. If this is an office setting, perhaps 10 hours is plenty. Just do the math and change the number of seconds to satisfy your lease time preference.
My DHCP Gateway is set to automatic. If my Gateway were something different, I could change it, but I'm using the default here.
I also have some various Advanced DHCP Options but there's nothing here that we need to configure. Just be aware that you do have some additional options if your organization happens to need them.
I'll click on Save here and go to the next step.
Now that we have our network, we now need to create a wireless network. To do that, I go to Wireless Networks.
I have a regular wireless network called CorpNet but I don't want this being used for personal devices. So I'll create a new one for that. I click on Create Wireless Network.
The first thing is to give it a Name or SSID. For that I just name it BYOD to stay consistent. I'll leave it Enabled.
By default it's set to be an Open network This is probably not a good idea since the traffic won't be encrypted and others on that network could easily have their traffic captured. I'll change this to WPA Personal and put in a security key here that my BYOD guest will use.
I'll check Guest Policy. This will enable extra security options such as dropping broadcast traffic from users' wireless connections. It's recommended that you check this.
Down here I said that we were using VLAN 50. This is where we put that in.
I'm not going to change any of the other settings. We have Rate and Beacon Controls, MAC Filtering, and we can configure RADIUS MAC Authentication. Once again, we won't configure any of this.
I'll click on Save.
So here you can see I have my BYOD wireless network in addition to my other wireless network.
Now let's look at Guest Control. I'll click on it and here I can see our settings.
First thing is to check the box to Enable Guest Portal.
Here I can decide what sort of authentication I might want to use. It's set to No authentication. I'll leave it set to the default.
I can choose how long someone has access before it times out. 8 hours is set by default but I can change that if I want to.
We can use the default landing page our BYOD guest will go to or have it redirected for them.

We have other settings such as creating a welcome message or displaying Terms of Service for those using our network.
We can customize the portal. This is what the BYOD guest sees when they connect. I can preview what the mobile users see as well, such as what they see on a notebook or desktop system. I can create a custom background image or custom logo and upload them instead of the default ones provided. I have all sorts of things I can change, like text, color, and so on.
For Access Control, I can decide what users have access to as soon as they connect as well as after they connect. I currently have three default networks that they're not able to access from the Guest network. I can add other networks if needed, or actually allow them access. I won't make any changes to this.
I'll click on Apply Changes.
The last thing I want to look at is User Groups. I'll start by clicking on Create New User Group. For name I'll put in BYOD. Now what we're doing here is limiting our users to the amount of bandwidth they're allowed to use. I'm sure you've been to public places and done a speed test only to find that the speed is horrible. Well, this is where the person responsible does this to you. Being greedy with my bandwidth, I'm only going to give my BYOD guests 5 Mbps for download. For upload I'm only going to give them 2 Mbps.
I'll click save and you can see the results. By the way, the default is unlimited bandwidth. That's whatever the max is. That's it for this demonstration.Â In this demo, we createdÂ a guest wireless network where users who bring their own devices can connect to the internet.

10.7.6 Create a Guest Network for BYOD (Simulation)

Scenario

You are a network technician for a small corporate network. You need to enable BYOD Guest Access Services on your network for guests and employees with mobile phones, tablets, and personal computers.

In this lab, your task is to perform the following:

Access the Wireless Controller console through Google Chrome at **http://192.168.0.6** .

Username: **admin** (case sensitive)

password: **password**

Set up Guest Access Services using the following parameters:

Name: **Guest_BYOD**

Authentication: **Use guest pass authentication**

The guest should be presented with your terms of use statement and allowed to go to the URL they were trying to access.

Verify that **192.168.0.0/16** is on the list of restricted subnets.

Create a guest WLAN using the following parameters:

Network name: **Guest**

ESSID: **Guest_BYOD**

Type: **Guest Access**

Authentication: **Open**

Encryption Method: **None**

Guest Access Service: **Guest_BYOD**

Isolate guest wireless clients from other clients on the access point.

Open a new Google Chrome window and request a guest pass using the BYODAdmin user as follows:

URL: **192.168.0.6/guestpass**

Username: **BYODAdmin** (case sensitive)

Password: **P@ssw0rd** (0 is a zero)

Use any ***full name*** in the Full Name field.

Make a note of or copy and paste the key in the Key field.

If you are using the Firefox browser, copying and pasting the key will not work. Make a note of the key value and type it into the Key field.

Use the key from the guest pass request to authenticate to the wireless LAN Guest_BYOD from the Gst-Lap laptop computer in the Lobby.

Explanation

Complete this lab as follows:

Access and log into the Ruckus **ZoneDirector** .

From the taskbar, select **Google Chrome** .

In the URL field, enter **192.168.0.6** and then press **Enter** .

Maximize the window for better viewing.

In the Admin field, enter **admin** (case sensitive).

In the Password field, enter **password** as the password.

Select **Login** .

Set up Guest Access Services.

Select the **Configure** tab.

From the left menu, select **Guest Access** .

Under Guest Access Service, select **Create New** .

Change the Name field to **Guest_BYOD** .

For Terms of Use, select **Show terms of use** .

Expand **Restricted Subnet Access** .

Verify that **192.168.0.0/16** is listed.

Select **OK** .

Create a Guest WLAN.

From the left menu, select **WLANs** .

Under WLANs, select **Create New** .

Change the Name to **Guest** .

Change the ESSID to **Guest_BYOD** .

Under Type, select **Guest Access** .

For Wireless Client Isolation, select **Isolate wireless client traffic from other clients on the same AP** .

Select **OK** .

Close **Google Chrome** .

Request a Guest password.

Open a new **Google Chrome** browser window.

In the URL field, enter **192.168.0.6/guestpass** and then press **Enter** .

Maximize the window for better viewing.

In the Username field, enter **BYODAdmin** (case sensitive).

Enter **P@ssw0rd** as the password (0 is a zero).

Select **Log In** .

In the Full Name field, enter any ***full name*** .

In the Key field, highlight the ***key*** and press **Ctrl + C** to copy the key.

If you are using the Firefox browser, copying and pasting the key will not work. Make a note of the key value and type it into the Key field.

Select **Next** .

Access the wireless Guest Access Service from the guest laptop in the lobby.

From the top left, select **Floor 1** .

Under Lobby, select **Gst-Lap** in the lobby.

In the notification area, select the wireless **Network** icon.

Select **Guest_BYOD** .

Select **Connect** .

Select **Yes** .

After Chrome opens to the Guest Access login page, paste the ***key*** from the Key field.

Select **Log In** .

10.7.7 Practice Questions (Section Quiz)

q_boyd_sec_aup_01_secp8

Which of the following is a policy that defines appropriate and inappropriate usage of company resources, assets, and communications?

Answers:

***Acceptable use policy (AUP)**

Business impact analysis (BIA)

Disaster recovery plan (DRP)

Business continuity plan (BCP)

Explanation:

An acceptable use policy (AUP) is a policy that defines appropriate and inappropriate usage of company resources, assets, and communications.

A business impact analysis (BIA) identifies critical processes and assets and the effect of their loss on the company.

A disaster recovery plan (DRP) addresses how a corporation should respond to a disaster.

A business continuity plan (BCP) addresses how a corporation responds to the disruption of critical systems.

q_boyd_sec_aup_02_secp8

Which of the following BEST describes an acceptable use agreement?

Answers:

***An agreement that identifies employees' rights to use company property, such as internet access and computer equipment, for personal use.**

An agreement that prohibits an employee from working for a competing organization for a specified period of time after he or she leaves the organization.

An agreement that outlines the organization's monitoring activities.

A legal contract between the organization and the employee that specifies that the employee is not to disclose the organization's confidential information.

Explanation:

An acceptable use agreement identifies employees' rights to use company property, such as internet access and computer equipment, for personal use.

A non-compete agreement prohibits an employee from working for a competing organization for a specified period of time after he or she leaves the organization.

An employee monitoring agreement outlines the organization's monitoring activities.

A non-disclosure agreement is a legal contract between an organization and an employee that specifies that the employee is not to disclose the organization's confidential information.

q_boyd_sec_aup_03_secp8

Your organization allows employees to bring their own devices into work, but management is concerned that a malicious internal user could use a mobile device to conduct an insider attack.

Which of the following should be implemented to help mitigate this threat?

Answers:

Implement a network access control (NAC) solution.

Implement an AUP that specifies which apps are allowed for use with organizational data.

Implement a guest wireless network that is isolated from your organization's production network.

***Implement an AUP that specifies where and when mobile devices can be possessed within the organization.**

Explanation:

To mitigate the threat of an insider attack, you should consider implementing an AUP that:

Specifies where and when mobile devices can be possessed within the organization. For example, the possession of mobile devices may be prohibited in high-security areas.

Notifies users that personally owned devices are subject to random searches if brought on site.

A network access control (NAC) solution would not help mitigate an insider attack with mobile devices.

Implementing an acceptable use policy (AUP) that specifies which apps are allowed for use with organizational data would not help mitigate an insider attack with mobile devices.

Implementing a guest wireless network that is isolated from your organization's production network would not help mitigate an insider attack with mobile devices.

q_boyd_sec_byod_secp8

Which of the following could be an example of a malicious insider attack?

Answers:

A user has lost a company-owned device.

A user's device has become infected with malware.

A user has not implemented appropriate security settings.

***A user uses the built-in microphone to record conversations.**

Explanation:

If a user is so inclined, he or she could use their mobile device to conduct a malicious insider attack. For example, they could:

Use the built-in camera, which nearly all modern mobile devices have, to take pictures of sensitive internal information.

Use the built-in microphone to record conversations.

Use the built-in video function to record proprietary processes and procedures.

Use the device's mobile broadband connection to transfer stolen data to parties outside the organization, bypassing the organization's network security mechanisms.

If a user copies sensitive data to their device, the organization could potentially lose control of that information. Even the question of who owns the data after it has been copied to a personal device becomes problematic. Consider the following scenarios:

A user may not have implemented appropriate security settings on their device, allowing anyone who gains access to the device to view sensitive data.

A user may lose the device, allowing anyone who finds it to access sensitive data.

A device may become infected with malware, potentially exposing sensitive data.

q_boyd_sec_cope_secp8

Which device deployment model gives businesses significant control over device security while allowing employees to use their devices to access both corporate and personal data?

Answers:

BYOD

CYOD

***COPE**

VDI

Explanation:

The corporate-owned personally enabled (COPE) model gives businesses significant control over device security while allowing employees to use their devices to access both corporate and personal data. Because the company owns the device, it can be secured more easily and wiped clean if lost or stolen. One disadvantage of this model is that employees who are not free to choose their own devices may end up bringing their own anyway.

The bring your own device (BYOD) model has users bringing in their personal devices and using them for business use.

The choose your own device (CYOD) model provides slightly more flexibility in giving users a limited selection of devices to choose from.

A virtual desktop interface (VDI) can be used with any device deployment model. A VDI allows mobile devices to establish a remote connection to a virtualized desktop.

q_boyd_sec_device_disconnect_secp8

You are the IT Security Manager at a mid-sized company. The company has recently adopted a bring your own device (BYOD) policy.

One of your employees, John, has been using his personal device for work-related tasks. John's device gets infected with malware, which he unknowingly connects to the company's network. As a result, the malware spreads across the network, compromising sensitive data.

As the IT Security Manager, what should be your immediate course of action?

Answers:

Ignore the issue as it's John's personal device and not the company's responsibility.

***Disconnect John's device from the network, isolate the affected systems, and start an investigation to understand the extent of the damage.**

Immediately fire John for causing a security breach.

Publicly blame John for the incident to set an example for other employees.

Explanation:

Disconnecting John's device from the network is the most appropriate response. Disconnecting John's device will prevent further spread of the malware. Isolating the affected systems will help contain the malware and prevent it from infecting other systems. Starting an investigation will help you understand the extent of the damage and how the malware was able to infect the network, which can inform future prevention strategies.

Ignoring the issue will not solve the problem. The malware could continue to spread and compromise more systems and data. As the IT Security Manager, it's your responsibility to ensure the security of the company's network and data, regardless of the source of the threat.

While John's actions led to a security breach, firing him immediately is not the best course of action. It's important to investigate the incident thoroughly to understand how it happened. It's also crucial to provide training to employees about the risks and responsibilities associated with the BYOD policy to prevent similar incidents in the future.

Publicly blaming John could create a negative work environment and may not necessarily prevent future incidents. Instead, use this incident as a learning opportunity for all employees. Provide training on the proper use of personal devices for work-related tasks and the potential security risks associated with them.

q_boyd_sec_mdm_secp8

Users in the sales department perform many of their daily tasks, such as emailing and creating sales presentations, on company-owned tablets. These tablets contain sensitive information. If one of these tablets is lost or stolen, this information could end up in the wrong hands.

The chief information officer wants you to implement a solution that can be used to keep sensitive information from getting into the wrong hands if a device is lost or stolen.

Which of the following should you implement?

Answers:

***A mobile device management (MDM) infrastructure**

A network access control (NAC) solution

An acceptable use policy (AUP)

A guest wireless network that is isolated from your organization's production network

Explanation:

A mobile device management (MDM) infrastructure, such as Microsoft Intune, can be used to wipe data clean from a device that has been lost or stolen.

A network access control (NAC) solution can remediate devices before allowing them to connect to your network.

An acceptable use policy (AUP) can be used to define which kind of data is allowed on personally owned devices and which kind of data is prohibited.

A guest wireless network that is isolated from your organization's production network allows user-owned devices to gain internet access, but it quarantines them from the rest of your organization's production network.

q_boyd_sec_nac_01_secp8

If a user's BYOD device (such as a tablet or phone) is infected with malware, that malware can be spread if that user connects to your organization's network. One way to prevent this event is to use a network access control (NAC) system.

How does an NAC protect your network from being infected by a BYOD device?

Answers:

***The NAC remediates devices before allowing them to connect to your network.**

The NAC forces BYOD devices to connect to a guest network that is isolated from your production network.

The NAC specifies which apps can be used while the BYOD device is connected to the organization's network.

The NAC notifies users that personally owned devices are subject to random searches if brought on site.

Explanation:

The NAC remediates devices before allowing them to connect to your network. This means that the NAC performs the following types of device management tasks before allowing a device to connect to the network:

Operating system updates

App updates

Anti-malware installation

Anti-malware definition updates

An alternative to using an NAC solution is to force BYOD devices to connect to a guest network that is isolated from your production network.

An acceptable use policy (AUP) specifies which apps can be used while the BYOD device is connected to the organization's network.

An AUP also notifies users that personally owned devices are subject to random searches if brought on site.

q_boyd_sec_nac_02_secp8

The IT manager has tasked you with implementing a solution that ensures that mobile devices are up to date, have anti-malware installed, and have the latest definition updates before being allowed to connect to the network.

Which of the following should you implement?

Answers:

BYOD

MDM

***NAC**

VDI

Explanation:

A network access control (NAC) solution can remediate devices before allowing them to connect to your network. This includes defining that a device is fully updated, has anti-malware installed, and has the latest definition updates.

The bring your own device (BYOD) model has users bringing in their personal devices and using them for business use.

A mobile device management (MDM) infrastructure, such as Microsoft Intune, can be used to track, manage, and even remotely wipe a user's mobile device.

A virtual desktop infrastructure (VDI) can be used with any device-deployment model. A VDI allows mobile devices to establish a remote connection to a virtualized desktop.

q_boyd_sec_wifi_secp8

Users in the sales department perform many of their daily tasks, such as emailing and creating sales presentations, on their personal tablets.

The chief information officer worries that one of these users might also use their tablet to steal sensitive information from the organization's network. Your job is to implement a solution that prevents insiders from accessing sensitive information stored on the organization's network from their personal devices while still giving them access to the internet.

Which of the following should you implement?

Answers:

A mobile device management (MDM) infrastructure

A network access control (NAC) solution

An acceptable use policy (AUP)

***A guest wireless network that is isolated from your organization's production network**

Explanation:

A guest wireless network that is isolated from your organization's production network allows user-owned devices to gain internet access, but it quarantines them from sensitive information on your organization's production network.

A mobile device management (MDM) infrastructure, such as Microsoft Intune, can be used to wipe data from a device that has been lost or stolen.

A network access control (NAC) solution can remediate devices before allowing them to connect to your network.

An acceptable use policy (AUP) can be used to define which kind of data is allowed and prohibited on personally owned devices.

q_boyd_vs_cyod_secp8

At a large tech company, the IT department explored options to accommodate employees who prefer using their own devices for work purposes. The management understands that embracing such a policy can improve productivity and job satisfaction.

They consider two strategies: bring your own device (BYOD) and choose your own device. The IT team evaluated the benefits and risks associated with each approach.

Which of the following statements about BYOD and CYOD are correct? (Select two.)

Answers:

***Bring your own device allows employees to use their personal devices for work purposes.**

***Choose your own device limits device choices to a predefined list of approved options.**

Choose your own device provides employees with the freedom to select any device for work, regardless of security concerns.

Bring your own device and choose your own device have the same level of control and management by the IT department.

Bring your own device (BYOD) allows businesses to purchase devices at significant discounts.

Explanation:

Bring your own device (BYOD) policy permits employees to utilize their devices, such as smartphones or laptops, for work tasks. This approach can increase employee satisfaction and flexibility but may pose security challenges if not properly managed.

Choose your own device (CYOD) allows employees to select their work devices from a preapproved list. This strategy offers a balance between employee choice and IT control.

CYOD allows employees to choose from a predetermined list of approved devices, ensuring better security and compliance with organizational policies.

BYOD involves managing personal devices that may vary in models and security configurations, whereas CYOD focuses on a limited set of approved devices, making management more structured and controlled.

Corporate-owned devices is an alternative to BYOD and CYOD. A corporate-owned device strategy lets businesses more effectively monitor and control activities performed on mobile devices. One advantage of this model is that businesses can purchase devices at significant discounts.

q_ipad_sec_secp7

Which version of iOS is currently running on your iPad?

Answers:

11.4

11.4.1

12.4

12.4.1

***15.2**

15.2.1

q_ipad_sec_secp7_02

What version of iOS is installed on your iPad after the update?

Answers:

11.2

11.2.1

12.4

12.4.1

15.2

***15.2.1**

10.8 Embedded and Specialized Systems

As you study this section, answer the following questions:

How can you minimize the damage of compromised embedded devices?

What are common static environments within the Internet of Things (IoT)?

In this section, you will learn to:

Configure smart home devices.

The key terms for this section include:

Term	Definition
Supervisory control and data acquisition (SCADA)	SCADA is an industrial computer system that monitors and controls a process.
Internet of Things (IoT)	The network of physical devices such as vehicles, home appliances, etc., that are embedded with electronics, software, sensors, actuators, and connectivity that enable them to connect, collect, and exchange data through the internet.
Arduino	Arduino is an open-source hardware and software platform for building electronic projects.
Raspberry Pi	Raspberry Pi is a low-cost device the size of a credit card that's powered by the Python programming language. It's manufactured into a single system on a chip (SoC).
Field Programmable Gate Array (FPGA)	FPGA (Field-Programmable Gate Array) is a reconfigurable integrated circuit that can be programmed to perform various tasks and functions.

Subscriber identity module (SIM) card	A SIM card encrypts data transmission and stores information.
Zigbee	Zigbee is a radio protocol that creates low-rate private area networks.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.4 Explain the importance of using appropriate cryptographic solutions. Tools Trusted Platform Module (TPM) 2.2 Explain common threat vectors and attack surfaces. Unsecure networks Wireless Wired Bluetooth 3.1 Compare and contrast security implications of different architecture models. Architecture and infrastructure concepts IoT Industrial control systems (ICS)/supervisory control and data acquisition (SCADA) Real-time operating system (RTOS) Embedded systems 3.3 Compare and contrast concepts and strategies to protect data. Methods to secure data Encryption

	4.1 Given a scenario, apply common security techniques to computing resources. Hardening targets Mobile devices ICS/SCADA Embedded systems IoT devices Mobile solutions Cellular Wi-Fi Bluetooth

10.8.1 Embedded and Specialized Systems (Lesson Video)

Transcript:

Today's smart technology is being embedded into devices like TVs, kitchen appliances, environmental controls, and even industrial equipment. And it's becoming increasingly common. These devices are sometimes referred to as static environments because they weren't designed to be customized or modified in any way by you, the network administrator.

Smart tech history is connected with supervisory control and data acquisition systems. We just call them SCADA. This infrastructure manages automated factory equipment that relies on the smart technology embedded in modern devices. As its name suggests, SCADA systems provide two general functions: supervisory control and data acquisition. The first controls remote equipment over a network connection. The second gathers information from those remote devices, allowing you to monitor their status.
Using network connectivity, SCADA systems can manage industrial equipment a long distance away from a control center. This is sometimes called a distributed control system, or DCS. To make this work, SCADA systems have several different components. First, there's a central supervisory computer that communicates with and sends control commands out to connected SCADA devices. The devices have remote terminal units, or RTUs. The RTUs connect the equipment to the network, which enables them to receive commands from the supervisory system and then send status information back.
Alternatively, SCADA equipment might have programmable logic controllers, or PLCs, installed. PLCs connect the SCADA equipment to the network just like RTUs. The difference is that they're usually less expensive.
Another important component is the type of network link used to connect the supervisory computer to the RTUs or PLCs. Some of the most common options include a standard internet connection, a satellite link, a private WAN link, or even a simple modem connection over a telephone line.
When talking about smart technology and SCADA, we can't forget the Internet of Things, or IoT.
IoT is a natural extension of SCADA. Its evolution is such that late-generation SCADA systems developed into first-generation IoT systems. An IoT ecosystem is made up of web-enabled smart devices that use embedded processors, sensors, and communication hardware to collect, send, and act on data they gather from their environments.
One important aspect of newer IoT devices is the operating system they use because these offer connectivity, usability, and interoperability.

One of these operating systems is called RIOT OS. It requires few resources and is very energy efficient. It's used on embedded systems, actuator boards, and sensors.
Another is ARM Mbed OS. This OS is primarily used with low-power technology such as wearable devices. In other words, any mobile device that's meant to be worn somewhere on the body. These
devices can be as simple as a step-tracking wristband or as complex as a virtual reality headset. Most wearable devices are designed to interface with another device.
For example, a smartwatch by itself has a limited set of functions. But when you connect it to a smartphone through Bluetooth, you get added functionality like text messaging or being able to answer phone calls.
Back to the operating systems now. RealSense OS X is used in Intel's depth-sensing technology for cameras and sensors. Then there's Nucleus and Integrity RTOS, which are both used for aerospace, industrial, automotive, and medical devices. When it comes to drones and robots, Ubuntu Core or Snappy are usually preferred.
Whether it's a tablet, smartphone, or an e-reader, all mobile devices share some characteristics. The most common is the power source. The power source is typically a lithium ion battery rated in milliamp hours. The larger the number, the greater the battery capacity. For example, the average smartphone battery is about 2,000 milliamp hours while a notebook battery is about 6,000 milliamp hours.
Another characteristic is a mobile device's storage.
To extend battery life, mobile devices use some sort of non-volatile flash memory to store information. Because no moving parts need to be powered, flash memory is a more energy-efficient storage medium. In addition, a lot of mobile devices have expandable storage in the form of micro-SD cards. This allows mobile devices to increase their storage to upwards of 200 gigabytes.
Mobile devices need connectivity to the internet and to other mobile devices. To do this, they use a variety of wireless connection mediums. For example, the typical smartphone utilizes 802.11, Bluetooth, and cellular wireless technologies.
Mobile devices also have internal sensors that collect environmental data, usually with a GPS chip. They use these chips to calculate the device's location information, which you can use for navigation and for finding your device if you happen to lose it.
Another internal sensor is called an accelerometer. This is a sensor that detects movement on a single plane, typically the horizontal plane. It detects when a mobile device is turned sideways and tells the device to change the screen's orientation.
A third sensor is a gyroscope. It also detects movement, but instead of detecting movement on a single plane, it detects horizontal and vertical movements. If you've ever played a game on a smartphone or tablet that requires you to tilt the device, you've used a gyroscope.
A key issue with smart technology is that there's little to no ability to modify the technology within. As a result, it's very difficult to manage security. For example, you can't install anti-malware software on a smart refrigerator or on an HVAC system even though there are significant security risks associated with them.
For example, in early 2014 the first documented IoT attack occurred. This attack leveraged smart devices to conduct a rather pervasive, malicious email-phishing exploit.
The trouble is that because smart devices are becoming more commonplace, we tend to forget about them. In addition, smart device vendors are slow to protect their products proactively because of cost constraints. For this reason, it's even more important for you as an administrator to secure your network against any possible smart device attacks.
It's also good to keep physical security in mind. When an attacker discovers a rogue access point, they're able to run various types of vulnerability scanners from outside the company. But if they gain physical access, they can hide a physical rogue access point as well. This is often done by configuring an extremely compact and powerful hardware device called a Raspberry Pi.
Just like smart device attacks, a rogue access point attack happens simply because a network administrator didn't properly configure the legitimate access points. Maybe he or she didn't implemented proper security measures in the beginning and eventually just forgot about them altogether. This is known as access point misconfiguration.
To defend against these attacks, there are a few things you can do. First, have a solid firewall design in place. Second, have intrusion detection. Or better yet, install intrusion-protection devices. Third, have very strong organizational security policies in place so that you're not caught off guard.
That's it for this lesson. In this video, we looked at security risks associated with embedded and specialized systems. We talked about SCADA, IoT, smart devices, and mobile devices. We talked about the security risks associated with these devices as well as with Raspberry Pi. We ended this lesson by talking about how you can defend yourself from distributed attacks leveled against these smart devices.

10.8.2 Smart Home (Demo Video)

Transcript:

Everyone seeks convenience and connectivity whether you are at work, home, in the car or wherever. Smart devices are showing up nearly everywhere. It could be your thermostat, light switches, light bulbs, electrical plugs, security cameras, door locks, sprinkler systems, and even digital speakers such as Google Home or Amazon Echo.

There are hundreds of these devices and more coming on the market all the time. It would be simply impossible to cover all of them, but we are going to take a look at a few of them in this demo.
The first one, and one of the very popular, is the Amazon Echo. The Echo is a smart speaker that is controlled by voice. It responds to voice commands when you use a "wake word," which is normally Alexa.
The device is capable of doing things such as playing music, making a to-do list, setting alarms, providing weather info, along with all sorts of other things. It can also be used in conjunction with other smart devices to control devices such as lights, security cameras, and so on. One of my favorite things to do at home is to use Bluetooth to pair Alexa with my surround sound system and stream music.
Most smart devices are managed with an app and the Echo is no different. Let's take a look at the Amazon Echo app on an Android Device. Here is a Google Pixel phone and I have my Amazon Alexa app on it. Let's go ahead and open it.
The first thing I can see are some things to try, followed by some of the recent things that it has done. As you can see, we have some music lovers that use the Echo.
Let's go up to our menu and then down at the bottom. We will tap on our settings. When we do that, we get a new screen. The first thing we can look at is the Alexa account. Here we do things such as train Alexa to recognize our voice. Perhaps we have an accent and want to train Alexa to recognize it better.
Below that, we can tell Alexa that it is OK to use voice commands to purchase items. If we go into that, we can see that we can require a code confirming the purchase of items. If that was not the case, I could see some issues with impulse shopping. Let's go back.
Next is the Alexa Voice Responses. Alexa is configured to respond in short replies which is the default.
The Amazon Household is a feature to allow other users to use this device to do things such as purchase products and services, using the Voice Purchasing feature.
Below that is History. Here you can see a list of things that Alexa has been asked to do. As you can see it can be quite detailed. Although you must use the "wake word" first, some feel that these smart speakers, may be listening to us a little too much at times.
The last selection will take you to the Amazon website where you can view more history items, frequently asked questions, the privacy notice, and the terms of use for the device.
The Echo can work with other smart devices. Let's tap down here on the lower right on the Smart Devices Icon.
The first thing we have is the Echo & Alexa group. If I tap on that, you can see that I have a couple of Echo's connected. If I tap on the first one, Dana's Echo, you can then see some addition information such as connected devices, the Wi-Fi information, and so on. Let's go back a couple of screens.
Next, we have Lights. This is a smart light located on the front door of my home. If I look at it, I can see that currently that light is on. If I tap it, it will go off. If I tap the icon below, Create a Routine, I could configure this light to come on at night and go off during the day. Now for this to work with this Echo, I had to use the Set this up on the app that came with the light, in order for this to work.
I'm going to go back a few screens and tap on Plugs. This is for a smart plug that is in our home and connected to a lamp. When I tap this here, I can turn this lamp off or on. This is also configured to work with voice commands. I can say, "Alexa, light off." And she would turn that light off.
If I tap on the Lamp, I can once again see that at the bottom I can tap here to create a routine for this plug to come off or on, either at certain times or use specially configured voice commands to do so.
Now, we saw how the smart plug was connected and we had a lamp plugged into it. Let's take a look at that app and see how that is configured. I'll exit my Amazon Echo app and go into my Smart Plug app.
OK, we are in our app and here you can see that I have one smart plug and I named it Lamp. From this app, I can turn the plug, which the lamp is plugged into, off and on. The button is green letting me know that it is now on.
If I had other smart devices from this manufacturer, they would be listed here. To add a device, I would tap the plus icon up here, and here we can see a list of all the different types of devices that I would be able to add. I don't have any additional devices but if I did, I would tap the type of device and this app would walk me through the installation and setup.
For example, if I had a smart bulb, I could just tap here. You can see that it tells me what to do. I don't have any of these, so I'll back out of here.

Finally, if I tap on the selection, Works with Kasa, I can see some of the other smart devices that are compatible with this device. You can see Amazon Alexa in this list. As we saw earlier, this smart plug works with my Amazon Echo.
The final thing we will look at are some security cameras. Once again, the manufacture uses a mobile device app to setup and configure the security cameras. The app is also used to view the video clips.
Now this system has three cameras set up and configured. Each one can be configured individually by tapping these setting next to the main icon. Here you can name each camera and check it's battery life. This system uses batteries, so they can be installed anywhere, as long as it can get my Wi-Fi signal.
It also shows me the temperature where the camera is mounted and if the camera is enabled or not. The retrigger time is the amount of time between the time it first senses motion and when I want it to start recording again a second or third time. I have this set to 10 seconds.
Since the street is in view of the camera and we have cars going by occasionally, I don't want it picking up every single car, so I have the sensitivity set to 3. Only large vehicles from the street set off this camera, but it still picks up anything that comes in the front yard.
The clip length is set to 25 seconds. If there is more motion after that, the camera will retrigger and record for another 25 seconds. However, below that you can see that the clip will end early if the motion stops.
As I scroll down, you can see other settings, Infrared Illuminator, video quality, if the camera is enabled, latest update, the Wi-Fi signal to the camera, and the camera to the sync module, which is located next to the wireless router. Let's go back a screen.
With this particular system, this top view is an image of what I'm looking at, not a live view. Since these run on battery, they are not always recording video, only on motion or if I tap the video camera icon over here on the left. When I do, it takes me to a live video stream of the camera. You can see that it's at night and I have a skeleton riding a bike next to my house.
OK, that's it for this demo. In this demonstration, we talked about smart devices. We looked at few common smart home devices. First, we looked at the Amazon Echo, then we looked at a smart plug and how that is able to work with the Echo. We ended with a quick overview of a home security camera system.

10.8.3 Constraints and Security of Embedded Devices (Lesson Video)

Transcript:

Embedded systems are usually dedicated to performing a precise function. They can vary in size, complexity, and function and can range from tiny, portable devices like smart watches to large traffic controllers. Some have low complexity, like a single micro-controller chip that opens and closes a gate, and some have very high complexity, like the embedded systems in an automated aircraft.

Nowadays, devices are becoming more intelligent and capable than ever. Manufacturers are always competing to be the first to release the latest and greatest device. But in the process of beating their competitors to market, many end up neglecting some of the most important security issues. These circumstances make it so devices are often sent out unchecked and undertested, making them easy for hackers to exploit.
In this lesson, I'll go over some of the security constraints that affect modern embedded devices.
Smart devices collect data for many reasons, such as improving efficiency, experience, and decision-making. Again, the problem is that some devices lack the most basic security and data-protection policies. Because you're dealing with sensitive information, you absolutely have to store and process data securely across your network. This means that you should redact sensitive data before you store it. An alternative is to use data separation to decouple personally identifiable information from data payloads. Data that isn't useful anymore should be disposed of in a safe manner. No measure is ever fail-proof, and every device needs regular patches to address attacks that uncover new flaws.
Patching has its downsides as well, and many times a patch ends up causing just as many problems as it was intended to solve. First of all, the patch might break the device's functionality. Second, most manufacturers are more concerned with producing and releasing products faster than their competition than mitigating security risks. Third, manufacturers often stop releasing patches and support for devices once they start working on a different product. Fourth, many devices use unsupported legacy Linux kernels.
Ideally, each device should undergo proper testing before it's released to the market, and patches should happen on a regular basis. But, sometimes the cost to go through all the testing is more than manufacturers are willing to spend. Many times, they simply accept the risks, even though skipping these important steps is bad for customers and

manufacturers. After all, it only takes one large-scale data breach for a company to lose sensitive information, the public's trust, and maybe their entire fortune.
For this reason, you should always implement secure authentication for users and apps. Authentication is an area that's often overlooked in embedded-technology development. The Mirai botnet was one of the biggest and most troublesome DDoS attacks in recent history. This attack was so successful because it found vulnerable IoT devices and used default usernames and passwords to log in and infect them.
Although almost all devices have to authenticate before they access a network, many allow easy-to-guess default passwords. The main problem is that there's no set regulation for smart devices that companies are legally forced to abide by. With this in mind, some ways to strengthen your own authentication practices are to use two-factor authentication, enforce strong passwords, and use certificates.
Even if you secure your devices with authentication, there's another vulnerability called lack of encryption. Data encryption and decryption is an ongoing process. Most network sensors aren't capable of supporting this process, which means that most data is transferred and received as cleartext. In addition, ports listening to the internet are many times left open continuously. Both factors make data extremely vulnerable to theft, breaches, and other malicious acts.
In order to try to standardize platforms and create trust in the embedded world, a group of experts got together in 2003 and created the Trusted Computing Group, or TCG. Together, they released a standard called Trusted Platform Module, or TPM. Although trust and security are often spoken of as one and the same, they're actually quite different.
Consider the fact that a secure system that can be bypassed without detection isn't trustworthy. At the same time, a trusted system with undiscovered vulnerabilities isn't secure.
To be trusted, a system must behave as expected and be able to detect that something is wrong. That's why TPM is so valuable to trusted computing. It provides essential services such as strong identity, secure storage, and integrity measurement.
That's it for this video. As embedded devices grow and develop, they become more susceptible to exploitation from hackers and other bad actors. Understanding the present challenges in the world of embedded devices can help you protect your data and stand out in your technological career.

10.8.4 Communication of Embedded Systems (Lesson Video)

Transcript:

Now that we have a better understanding of embedded systems, it's time to dive deeper into the technologies and protocols that allow them to communicate.

Embedded systems often require new technologies, and their biggest obstacle is the lack of capable support. For this reason, they currently use standard networking protocols. These protocols are divided into four categories: short-range, medium-range, long-range, and wired.
There are several short-range protocols and technologies. Near-field communication, or NFC, and radio-frequency identification, or RFID, are two very simple, low-energy, and versatile protocols. RFID uses two-way radio transmitter receivers to identify and track object tags. NFC uses magnetic field induction to communicate between mobile and standard electronic devices.
Bluetooth Low Energy, also known as BLE or Bluetooth Smart, is a wireless personal area network. It supports low-power, long-use IoT technology. It's used in the healthcare, security, entertainment, and fitness sectors.
Light-Fidelity, or Li-Fi, is very similar to Wi-Fi. The two key differences are speed and mode of communication. Unlike Wi-Fi, Li-Fi is a visible light communications system. It uses light bulbs to transfer data at the high speed of 224 gigabits per second.
Quick Response, or QR, codes and barcodes are tags attached to products. They're machine-readable and contain information about the product. A QR code is two-dimensional, and you can scan it using a smartphone, while barcodes come in one-dimensional or two-dimensional code.
Thread, Zigbee, and Z-Wave are all radio protocols that create low-rate private area networks. Their advantage is that, although they're low-power, they offer high throughput. Thread uses an IPv6-based networking protocol. Zigbee is a short-range communication protocol based on the IEEE 203.15.4 standard. Z-Wave is a low-power, short-range communication protocol designed mostly for home IoT systems.
Wi-Fi is commonly implemented in wireless local area networking. The most common Wi-Fi standard is 802.11n with a maximum speed of 600 Megabits per second and a range of about 50 meters.

Wi-Fi Direct uses peer-to-peer communication without a set wireless access point. The devices in Wi-Fi Direct only start communicating after an access point device has been selected within the system.
You can also use hotspots to provide wireless network access.
Mobile hotspots let you connect a device to the internet through a portable wireless device such as a phone.
Hotspots form an on-the-spot Wi-Fi network, allowing you to connect a number of gadgets for simple, fast internet access. Hotspots often use third-, fourth-, and fifth-generation technology to provide this type of connection. 5G is the best of them because it achieves speeds twenty times faster that 4G; its peak speed is 20 gigabytes per second. It uses LTE for wireless connections and includes lower frequencies than previous generations, down to 600 Megahertz.
There are two main protocols in the medium-range category, HaLow and LTE-Advanced. HaLow is a branch of Wi-Fi with extended range. It's most useful in rural areas because it uses low data rates, so it reduces power and cost for transmission. LTE-Advanced is a mobile communication. It makes standard LTE better because it has a greater capacity for data rate, extended range, efficiency, and performance.
Long-range wireless communication offers a few different protocols and technologies. The first one we'll talk about is LPWAN. LPWAN stands for low-power wide-area networking. It's a wireless telecommunication network. There are three protocols associated with LPWAN: LoRaWAN, Sigfox, and Neul.
LoRaWAN stands for long-range wide-area network. It's used with mobile, industrial machine-to-machine, and secure two-way communication for IoT devices, smart cities, and healthcare. Sigfox is good for devices with short battery life that need to transfer low-level data. Neul uses a small part of the TV white space spectrum and delivers high-quality, high-power, high-coverage networks at a low cost. Next, there's Very Small Aperture Terminal, or VSAT. VSAT uses small dish antennas to transfer broadband and narrowband data.
And finally, there's cellular. This communication protocol can send and receive high-quality data over very long distances. However, it's very expensive and consumes a lot of power.
Next we're going to talk about wired communication. There are three main protocols in this category: Ethernet; Multimedia over Coax Alliance, or MoCA; and Power-line Communication, or PLC.
Ethernet is used in LANs. The speed, signaling, media access methods, and cable lengths are defined in the various Ethernet standards published by the IEEE 802.3 committee. There are many Ethernet specifications, but they all use a common naming standard. The word BASE is in the middle. This indicates that the standard uses baseband signaling, not broadband signaling. Baseband updates may improve or diminish battery performance, network signal strength, and roaming capability.
Then we have MoCA. As its name suggests, MoCA uses coaxial cables to provide high-definition videos of a home and other content related to it.
PLC uses electrical wires to transmit power and data from one point to another. PLC is used in different sectors like automation, industrial, and broadband over power lines, or BLP.
One last form of embedded device communication worth mentioning isn't a protocol; it's a card many devices carry inside them called a Subscriber Identity Module card, or SIM card. It encrypts data transmission and stores information about the owner of the device the SIM card is in for identification and authentication. It also stores the International Mobile Subscriber Identity, or IMSI. This is a unique identifier that defines a subscriber in the wireless world, including their country and mobile network.
That's it for this lesson. Today, we discussed the communication of embedded systems. We talked about short-range, medium-range, and long-range protocols. We also covered wired communication and SIM cards. Embedded devices are complex. Learning about their protocols, technologies, and communication models will prepare you to incorporate and safeguard them.

10.8.5 Embedded and Specialized Systems Facts

More and more devices are becoming connected to the internet through embedded technology that allows the device to send and receive information.

This lesson covers the following topics:

- Types of embedded devices
- Security risks

Types of Embedded Devices

The following table describes some of the most common embedded devices:

Device Type	Description
Appliances	Many appliances contain integrated technology that allows internet communication. For example, smart laundry appliances can send notifications when a load is complete or when laundry detergent needs to be refilled. Other common smart home appliances include: Refrigerators Dishwashers Microwaves
Environment controls	Many homes and businesses use environmental control devices that can send real-time information and can be controlled via the internet. These devices can be as basic as controlling a home's HVAC system (such as a Nest thermostat) or as complex as controlling the humidity, temperature, and other environmental factors in a data center.
Building/facility automation	Some facilities use a network of integrated devices that control various aspects, creating what is known as facility automation. Some of the devices that are integrated with facility automation include: Lighting controls Security systems Door locks Sprinkler systems Garage doors Smart meters
Wearable devices	In recent years, companies have started producing wearable devices that can connect to the internet for a variety of purposes. These devices include: Watches Headphones Fitness trackers
Automobiles	Modern cars use integrated technologies and in-vehicle systems that can perform various tasks, such as: Starting the car remotely using a smartphone.

	Warning a driver about nearby cars. Applying the brakes automatically to avoid a collision. Performing parallel parking autonomously.
Industrial equipment	Some industrial equipment also fits into the category of a smart device. Supervisory control and data acquisition (SCADA) devices are special computer systems that gather, analyze, and manage automated factory equipment. For example, a SCADA system could be used to monitor factory pipes and automatically open valves if pressure in the pipe system reaches a specific threshold. SCADA is a subset of Industrial Control Systems (ICS), which refers to all types of industrial automation.
Mainframe computer	A lesser-known category of embedded devices is mainframe computers. A mainframe computer is a large, powerful computer that is capable of processing extremely large amounts of data. Mainframe computers typically run proprietary operating systems. Because these operating systems are rarely updated, they are considered a static environment. In addition, mainframe computers often contain large amounts of sensitive data, making them an attractive target for hackers.
Real-time operating system (RTOS)	A RTOS is an operating system that serves real-time applications without buffer delays. They are generally used in systems that require a response within a strict time constraint. Because RTOS are often used as critical components of an application, a successful attack on the RTOS can harm an entire system, including physical machinery.
System on a chip (SoC)	A SoC is an integrated circuit that includes all components of a typical computer system, including digital, analog, mixed-signal, and radio frequency functions. The Raspberry Pi is a common device that uses a SoC. Because of their relatively low cost, SoCs are often used by hobbyists.
Multi-function display (MFD)	An MFD is a screen surrounded by configurable buttons that can be used to display information in a variety of ways. MFDs are often used on airplanes, helicopters, and ships.
Medical devices	Much of today's medical technology for daily monitoring and maintenance uses embedded systems. Instead of having to visit a physician every day, wearable devices can be used to collect information on heart rate, glucose levels, weight, blood pressure, and other parameters. This information can then be sent to a doctor automatically or used for self-monitoring.
Unmanned aerial vehicles (UAV)	UAVs are used for military campaigns, search and rescue, weather monitoring, and recreation. UAVs use embedded computers for collecting and transmitting data and for receiving commands.
Digital cameras	Most modern digital cameras use embedded systems for processing captured images, storing images, and uploading images to a PC or other storage device.
Media gateways	A media gateway is a translation device that converts media streams for use by different telecommunication technologies.

Wireless keyboards and mice	Wireless keyboards and mice use Bluetooth or other proprietary radio frequency connections.
Displays	In the past, display devices had a single use as a monitor for a computer. Today's monitors and other display devices are increasingly embedded with smart features and have wireless connections.
Wi-Fi-enabled microSD cards	Wi-Fi-enabled MicroSD cards can wirelessly transfer data to and from other devices. Many of them connect directly to the internet.
Multifunction printers (MFPs)	Multifunction printers can connect to wireless networks and the internet for additional functionality.
External storage devices	External storage devices such as USB flash drives, HDDs, and SSDs can connect to traditional computing equipment, as well as to many smart devices.
Arduino	Arduino is an open-source hardware and software company. They design and manufacture single-board microcontrollers as well as kits to build digital devices.
Field Programmable Gate Array (FPGA)	A Field Programmable Gate Array is an integrated circuit manufactured and then later configured by the customer. The configuration happens through a hardware description language (HDL), similar to an application-specific integrated circuit (ASIC).
Voice over IP (VoIP)	Voice over IP is a protocol optimized for the transmission of voice data (telephone calls) through a packet-switched IP network. VoIP routes phone calls through an IP network, including the internet. VoIP solutions can integrate with the public switched telephone network (PSTN) to allow VoIP customers to make and receive external calls.

Security Risks

As with any networked system, there are security risks associated with smart devices. Not only do you have little or no control over the smart technology within static environments, but smart device vendors can be slow to take steps to protect their products against security threats. They tend only to respond after an exploit has occurred instead of proactively updating systems. This is why smart devices are attractive to hackers. However, there are some steps you can take to secure a network from these devices and reduce the damage that a compromised device can cause.

- Some static devices (such as home routers, game consoles, and SCADA devices) require manual firmware updates. With these devices, it is important to keep the firmware updated.
- For devices that cannot be manually updated, the best approach is to minimize the amount of damage a compromised device can cause. This is done by segmenting the network using VLANs or encrypting all network communications.

10.8.6 Practice Questions (Section Quiz)

q_embed_sys_arduino_secp8

Which of the following is an open-source hardware and software company that designs and manufactures single-board microcontrollers as well as kits to build digital devices?

Answers:

Raspberry Pi

Microsoft

Amazon

***Arduino**

Explanation:

Arduino is an open-source hardware and software company. They design and manufacture single-board microcontrollers as well as kits to build digital devices.

Raspberry Pi is a common device that uses a system on a chip (SoC).

Neither Microsoft nor Amazon are an open-source hardware and software company that designs and manufactures single-board microcontrollers as well as kits to build digital devices.

q_embed_sys_facility_secp8

You manage information systems for a large co-location data center.

Networked environmental controls are used to manage the temperature within the data center. These controls use embedded smart technology that allows them to be managed over an internet connection using a mobile device app.

You are concerned about the security of these devices.

What can you do to increase their security posture? (Select two.)

Answers:

***Verify that your network's existing security infrastructure is working properly.**

Install anti-malware software on each device.

Enroll each device in a mobile device management (MDM) system.

Rely on the device manufacturer to maintain device security with automated firmware updates.

***Install the latest firmware updates from the device manufacturer.**

Explanation:

Since you generally have little or no control over the embedded technology within smart environmental control devices, they are referred to as static environments. As a result, there is typically very little you can do to increase the security posture for these types of devices. For environmental controls, you may be able to perform the following, depending upon the device manufacturer:

Install the latest firmware updates from the device manufacturer.

Verify that your network's existing security infrastructure is working properly.

Because these devices operate in a static environment, you typically can't install third-party software on them, including anti-malware scanners or mobile device management (MDM) agents.

Relying on the device manufacturer for security updates is problematic because manufacturers can be slow to take steps to protect their products against security threats.

Manufacturers tend only to respond after an exploit has occurred instead of proactively defending their systems.

q_embed_sys_fpga_secp8

Which of the following embedded devices is an integrated circuit that is manufactured and then later configured by the customer through a hardware description language (HDL)?

Answers:

System on a chip (SoC)

Multi-function display (MFD)

***Field-programmable gate array (FPGA)**

Real-time operating system (RTOS)

Explanation:

Field-programmable gate array (FPGA) is the correct answer. An FPGA is an integrated circuit that is manufactured and then later configured by the customer through a hardware description language (HDL).

A SoC is an integrated circuit that includes all components of a typical computer system, including digital, analog, mixed-signal, and radio frequency functions. However, it is not configured by the customer after manufacturing.

An MFD is a screen surrounded by configurable buttons that can be used to display information in a variety of ways. It is not an integrated circuit that is configured by the customer after manufacturing.

An RTOS is an operating system that serves real-time applications without buffer delays. It is not an integrated circuit that is configured by the customer after manufacturing.

q_embed_sys_manu_secp8

You manage the information systems for a large manufacturing firm.

Supervisory control and data acquisition (SCADA) devices are used on the manufacturing floor to manage your organization's automated factory equipment. The SCADA devices use embedded smart technology, allowing them to be managed using a mobile device app over an internet connection.

You are concerned about the security of these devices.

What can you do to increase their security posture? (Select two.)

Answers:

***Install the latest firmware updates from the device manufacturer.**

***Verify that your network's existing security infrastructure is working properly.**

Install anti-malware software on each device.

Enroll each device in a mobile device management system.

Install a network monitoring agent on each device.

Explanation:

Since you generally have little or no control over the smart technology embedded within SCADA devices, they are referred to as static environments. As a result, there is typically very little you can do to increase the security posture for these types of devices. For SCADA devices, you may be able to perform the following, depending on the device manufacturer:

Install the latest firmware updates from the device manufacturer.

Verify that your network's existing security infrastructure is working properly.

Because these devices operate in a static environment, you typically can't install third-party software on them, including anti-malware scanners, monitoring agents, or mobile device management agents.

q_embed_sys_rtos_01_secp8

Which of the following serves real-time applications without buffer delays?

Answers:

SCADA

SoC

FPGA

***RTOS**

Explanation:

A real-time operating system (RTOS) is an operating system that serves real-time applications without buffer delays. They are generally used in systems that require a response within a strict time constraint.

Supervisory control and data acquisition (SCADA) devices are special computer systems that gather, analyze, and manage automated factory equipment.

A system on a chip (SoC) is an integrated circuit that includes all components of a typical computer system, including digital, analog, mixed-signal, and radio frequency functions.

A field-programmable gate array (FPGA) is an integrated circuit manufactured and then later configured by the customer.

q_embed_sys_rtos_02_secp8

At a manufacturing facility, the security team for an IoT-based system implemented various efforts to monitor and control multiple processes using embedded systems.

The security team employed real-time operating systems (RTOS) to ensure precise and timely data processing. The security team recognized the importance of safeguarding these IoT devices and embedded systems against potential cyber threats.

The team plans to implement security BEST practices and protocols.

How would the security team explain why they want to implement security best practices and protocols in the manufacturing facility's IoT devices and embedded systems?

Answers:

***To reduce the risk of unauthorized access and data breaches in the RTOS and embedded systems.**

To ensure the precise monitoring and control of manufacturing processes.

To develop new security protocols tailored specifically for IoT devices.

To improve the manufacturing facility's overall efficiency and productivity.

Explanation:

The team wants to implement security best practices and protocols to reduce the risk of unauthorized access and data breaches in the real-time operating systems, Internet of Things (IoT) devices, and embedded systems used in the manufacturing facility.

Ensuring the precise monitoring and control of manufacturing processes is not the primary focus of security best practices and protocols. Instead, security best practices and protocols aim to enhance security rather than precise monitoring and control.

Developing new security protocols is irrelevant to implementing established best practice security protocols applicable to various devices, including IoT.

Efficiency and productivity do not relate to implementing security best practices and protocols. Instead, they are essential in business operations.

q_embed_sys_scada_secp8

Which of the following devices are special computer systems that gather, analyze, and manage automated factory equipment?

Answers:

***SCADA**

SoC

MFD

UAV

Explanation:

Supervisory control and data acquisition (SCADA) devices are special computer systems that gather, analyze, and manage automated factory equipment.

A system on a chip (SoC) is an integrated circuit that includes all components of a typical computer system, including digital, analog, mixed-signal, and radio frequency functions.

A multi-function display (MFD) is a screen surrounded by configurable buttons that can be used to display information in a variety of ways.

Unmanned aerial vehicles (UAVs) are used for military campaigns, search and rescue, weather monitoring, and recreation.

q_embed_sys_soc_secp8

Which of the following do Raspberry Pi systems make use of?

Answers:

SCADA

***SoC**

FPGA

RTOS

Explanation:

A system on a chip (SoC) is an integrated circuit that includes all components of a typical computer system, including digital, analog, mixed-signal, and radio frequency functions. Raspberry Pi is a common device that uses an SoC. Because of their relatively low cost, SoCs are often used by hobbyists.

A real-time operating system (RTOS) is an operating system that serves real-time applications without buffer delays. They are generally used in systems that require a response within a strict time constraint.

Supervisory control and data acquisition (SCADA) devices are special computer systems that gather, analyze, and manage automated factory equipment.

A field-programmable gate array (FPGA) is an integrated circuit configured by the customer.

q_embed_sys_uav_secp8

Which of the following embedded devices is primarily used for tasks such as military campaigns, search and rescue, weather monitoring, and recreation?

Answers:

Appliances

Environment controls

***Unmanned aerial vehicles (UAV)**

Multi-function display (MFD)

Explanation:

Unmanned aerial vehicles (UAV) is the correct answer. UAVs are used for a variety of tasks including military campaigns, search and rescue, weather monitoring, and recreation. They use embedded computers for collecting and transmitting data and for receiving commands.

While many modern appliances do contain integrated technology that allows for internet communication, they are not primarily used for tasks such as military campaigns, search and rescue, weather monitoring, or recreation.

Environment controls are devices that send real-time information and can be controlled via the internet. They are used for controlling a home's HVAC system or controlling the humidity, temperature, and other environmental factors in a data center, not for tasks such as military campaigns, search and rescue, weather monitoring, or recreation.

A multi-function display (MFD is a screen surrounded by configurable buttons that can be used to display information in a variety of ways. They are often used on airplanes, helicopters, and ships, but they are not used for tasks such as military campaigns, search and rescue, weather monitoring, or recreation.

q_embed_sys_vlan_secp8

As a network security analyst, you have been tasked with improving the security of a network that includes a variety of embedded devices, including appliances, wearable devices, and industrial equipment.

The network has been experiencing frequent security breaches.

Which of the following would be the MOST effective strategy to improve network security?

Answers:

Implementing a firewall for each individual device.

***Using VLANs or encrypting all network communications.**

Regularly updating the firmware of all devices.

Disabling all unnecessary services on the devices.

Explanation:

Using VLANs or encrypting all network communications is the correct answer. Using VLANs can help segment the network and isolate traffic, reducing the potential impact of a security breach. Encrypting all network communications can help protect data in transit and prevent unauthorized access.

While firewalls can provide an additional layer of security, implementing a firewall for each individual device can be resource-intensive and may not be feasible or effective in a network with a variety of embedded devices.

While regularly updating the firmware of all devices can help address known vulnerabilities, it may not be sufficient to protect against all potential security breaches, especially if the network communications are not secure.

While disabling unnecessary services can reduce potential attack vectors, it may not be sufficient to protect against all potential security breaches, especially if the network communications are not secure.

q_embed_sys_voip_secp8

Which of the following lets you make phone calls over a packet-switched network?

Answers:

***VoIP**

FPGA

RTOS

SCADA

Explanation:

Voice over IP (VoIP) is a protocol optimized for the transmission of voice data (telephone calls) through a packet-switched IP network. VoIP routes phone calls through an IP network, including the internet. VoIP solutions can integrate with a public-switched telephone network (PSTN) to allow VoIP customers to make and receive external calls.

A field-programmable gate array (FPGA) is an integrated circuit configured by the customer.

A real-time operating system (RTOS) is an operating system that serves real-time applications without buffer delays. They are generally used in systems that require a response within a strict time constraint.

Supervisory control and data acquisition (SCADA) devices are special computer systems that gather, analyze, and manage automated factory equipment.

10.9 Email

As you study this section, answer the following questions:

How does spam filtering help end users?

In what format are emails sent?

Why is it important to add multiple layers of security?

Why would you encrypt email coming only from outside your network?

What is S/MIME?

What is the difference between POP3 and IMAP?

In this section, you will learn to:

Protect a client from spam.

Secure an email server.

Configure email filters.

Secure accounts on an iPad.

Secure email on an iPad.

The key terms for this section include:

Term	Definition
Spam	Unwanted and unsolicited email usually sent to many recipients.
SMTP relay	An email server that accepts mail and forwards it to other mail servers.
Phishing email	A fraudulent email claiming to be from a trusted organization. The email typically asks a user to verify personal information or send money.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.4 Given a scenario, analyze indicators of malicious activity. Malware attacks Virus Application attacks 3.2 Given a scenario, apply security principles to secure enterprise infrastructure. Secure communication/access Tunneling Transport Layer Security (TLS) 4.1 Given a scenario, apply common security techniques to computing resources. Application security Secure cookies 4.5 Given a scenario, modify enterprise capabilities to enhance security. Email security Domain-based Message Authentication Reporting and Conformance (DMARC)

	DomainKeys Identified Mail (DKIM) Sender Policy Framework (SPF) Gateway DLP 5.6 Given a scenario, implement security awareness practices. Phishing
TestOut Security Pro	3.2 Implement Application Defenses 3.2.1 Implement an application allow list 3.2.4 Configure email filters and settings 3.2.5 Configure browser settings

10.9.1 Email Security (Lesson Video)

Transcript:

Email security is a big issue for IT professionals. Spam, viruses, malware, social engineering all use email as the vehicle to carry out an attack. In addition, a lot of sensitive information is sent via email.

In this lesson, we're going to talk about the steps you should take to properly secure email communications and reduce the risk of email exploits. Let's start by looking at a comprehensive email security solution called an email security gateway.
An email security gateway is a security solution that monitors emails that are sent to or originate from an organization.
An email security gateway can be software-, hardware-, or cloud-based. There are even some virtualized email security gateways that function on a hypervisor.
Hardware-based gateways are typically installed in a server room and are configured to have all email messages routed through the device. Software- or cloud-based gateways function in the same way by routing email messages through the service.
Depending on the solution, an email security gateway can offer things like spam protection, malware and virus scanning, email encryption, and even data loss prevention. Let's take a look at some of the features that most email security gateways offer.
Spam filtering is offered by a lot of gateways. As its name suggests, spam filtering is to filter out certain types of email. The filter will try to filter out unwanted or unsolicited emails and prevent them from arriving to the end users.
Spam filters use various methods in order to identify spam. Some services even use lists that are continually updated with known spammers and automatically block anything from those emails or domains.
One thing to know about spam filters is that they are far from perfectâ€"they are not a silver bullet. In fact, fine-tuning a spam filter is a never-ending process. Spammers are constantly finding ways to trick spam filters in order to get emails through and sometimes legitimate emails will be flagged and blocked by the filtering device. However, the benefits of using a spam filter far outweigh the time it takes to configure one.
Another feature provided by a lot of email security gateways is that of data loss prevention, or DLP.
Gateways that employ DLP scan all outgoing emails for sensitive information. Both the message body and attachments are scanned and if sensitive information is found, the email will be blocked. If desired, a notification can also be sent to the sender or system administrator.

For example, let's say an employeeâ€""either accidently or purposelyâ€""sends an email that contains a customer's credit card information to a personal email address. When the email is sent, it is filtered through the email gateway where it is scanned by the DLP feature. Because the email contains sensitive information, the gateway blocks the email from being sent outside the network. The gateway then notifies the employee that their email was blocked and that the system administrator has been notified.

Just like a spam filter, a DLP solution isn't prefect. Its success depends entirely on how advanced the scanning algorithm is and how well the gateway was configured. Because of this, it's important you understand how to properly configure DLP on a gateway.

The next email security gateway feature we'll look at is encryption.

Remember, email sent over the internet isn't very secure. Emails can be intercepted by someone with the right know how. And because email messages aren't encrypted by default, they can be easily read. To protect against this, you can use email encryption as a standalone or as part of an email security gateway.

As its name suggests, email encryption encrypts the entire contents of an email and, in some cases, will digitally sign the email with a certificate. This certificate is similar to SSL certificates in that they provide a form of identity verification.

Email encryption can be configured in a few different ways.

One way is to encrypt all email communications, regardless of destination. Doing this does provide the most protection; however, encrypting and decrypting emails requires a lot of resources. And if hundreds or even thousands of emails are being sent a day, this might not be feasible.

Another option is to only encrypt emails that are being sent outside the internal network. With this method, that funny cat video you email to Rick in accounting won't be encrypted, but the contract document you send to your client will be.

Now, even if you only encrypt outgoing emails, this encryption and decryption can still become a resource bottleneck. Luckily, some security gateways offer selective encryption. With this option, outgoing emails are scanned in a similar way that a DLP solution scans emails. If sensitive information is found, then the email is encrypted before its sent. If no sensitive information is found, then the email is sent unencrypted.

Now, this encryption and digital signing can be done using several different encryption protocols, such as PGP, GPG, or TLS. However, the most common and widely accepted protocol for email encryption is the S/MIME protocol, which stands for Secure/Multipurpose Internet Mail Extensions. The important thing to understand about S/MIME is that it requires the use of a public key certificate in order to encrypt and decrypt email messages. This certificate can either be self-signed or purchased from a public CA, but it is required.

Typically, two certificates are used: one for encrypting the email and another for digitally signing the email. Digitally signing the email ensures that only the intended recipient is able to decrypt the email message and also verifies that the email originated from a trusted source.

In addition to using an email security gateway, you can also further secure email communications by ensuring you are using the latest email protocols.

Whether you are using POP or IMAP, make sure to use the latest version of the protocol. In addition, both of these protocols support the use of SSL/TLS in order to create a secure tunnel when connecting to an email server. One way to do this is to use a Microsoft Exchange server and enable secure POP or secure IMAP communication.

Proper email security is an extremely important aspect of network security. Emails can be a vehicle for countless attacks and exploits. As such, you need to make sure you take the appropriate steps in order to properly secure your organizations email communications and prevent email-driven attacks.

10.9.2 Email Security Facts

This lesson covers the following topics:

- Email threats
- Email security

Email Threats

To secure email, you must be aware of the following email attacks.

Attack	Description
Virus	A virus is malware that often uses email as its distribution mechanism. Users receive the virus as an attachment and then activate the virus by clicking on that attachment. To mitigate viruses, you should install antivirus software on every system and install antivirus software on the email server to scan attachments. A best practice is to detect viruses and messages on the email server before they get to the client and send warnings to the recipient about the malicious email.
Spam	Spam is unwanted and unsolicited email sent to many recipients. Spam consists of the following attributes: Can be as benign as emails trying to sell products Can be malicious and contain phishing content, drive-by downloads, or malware Can contain malware as attachments Wastes bandwidth and could fill the inbox, resulting in a denial-of-service condition To control spam: Enable spam filters on the client and email servers. Filter junk email by identifying safe senders (whitelists), blocked senders (blacklists), countries to block email from, and languages to block. Enable antivirus scanning for attachments on the client and email servers. In the email client, disable preview screens. An email can have links for active items that can report back to the spammer. Do not click on an unsubscribe link at the bottom of an unsolicited email. Doing this verifies to the spammer that the email address is a current and active email address. Only unsubscribe from trusted organizations. Install server-level anti-spam software on the email server. Do not post your full email address anywhere on the web. Spammers use software to scan websites to find email addresses and then add them to their email lists for spamming.
Open SMTP relay	An SMTP relay is an email server that accepts mail and forwards it to other mail servers. An open SMTP relay allows anyone to forward mail. If your mail server is an open SMTP relay, it can be used by spammers to send mail. Spammers use your relay to obscure the actual source of the email. A repudiation attack is an attack on open relays in which the attacker accesses your email server and sends spoofed emails to others, making them appear as if they came from you.

	If spammers use your relay for sending mail, your server will soon be placed on a blacklist. Other mail servers then stop receiving any mail (even legitimate mail) sent from your servers. As a best practice: Configure your mail server to accept mail only from authenticated users or specific email servers that you authorize. Require TLS encryption to connect to the server. Implement restrictions for accessing the server and relaying email for your environment if feasible.
Phishing	A phishing email is an email pretending to be from a trusted organization that asks to verify personal information or send money. In a phishing attack: A fraudulent message (that appears to be legitimate) is sent to a target. The message requests that the target visit a fraudulent website (which also appears to be legitimate). Graphics, links, and web pages look almost identical to the legitimate requests and websites the attacker tries to represent. The fraudulent website requests that the victim provide sensitive information, such as the account number and password. To protect against phishing: Check the email header information to see more info about the sender and the links that are in the email. Only open emails if you recognize the sender. Check the actual link destination within emails to verify that they go to the correct URL and not a spoofed one. Do not click on links in emails. Instead, type the real URL into the browser. You could also look up the website in a search engine. Verify that HTTPS is used when going to e-commerce sites. HTTPS requires a certificate that matches the server name in the URL that is verified by a trusted CA. You can also look for the lock icon to verify that HTTPS is used. Implement phishing protections within your browser.

In addition, there are several protocols that help identify and mitigate phishing and spam emails by authenticating sender identities and checking message integrity, including protocols such as SPF (Sender Policy Framework), DKIM (DomainKeys Identified Mail), and DMARC (Domain-based Message Authentication, Reporting & Conformance). These protocols establish trust in email communications, reducing the likelihood of successful phishing and spam attacks.

Below are a few characteristics of the protocols mentioned above:

Protocol	Description
Sender Policy Framework (SPF)	This is an email authentication method that helps detect and prevent sender address forgery commonly used in phishing and spam emails. SPF works by verifying the sender's IP address against a list of authorized sending IP addresses published in the DNS TXT records of the email sender's domain. When an email is received, the receiving mail server checks the SPF record of the sender's domain to verify the email originated from one of the pre-authorized systems.
DomainKeys Identified Mail (DKIM)	This protocol leverages encryption features to enable email verification by allowing the sender to sign emails using a digital signature. The receiving email server uses a DKIM record in the sender's DNS record to verify the signature and the email's integrity.
Domain-based Message Authentication, Reporting and Conformance (DMARC)	DMARC uses the results of SPF and DKIM checks to define rules for handling messages, such as moving messages to quarantine or spam, rejecting them outright, or tagging the message. DMARC also provides reporting capabilities, giving the owner of a domain visibility into which systems are sending emails on their behalf, including unauthorized activity.

Email Security

Email is cleartext by default. To secure email, use either Secure/Multipurpose Internet Mail Extensions (S/MIME) or Pretty Good Privacy (PGP).

- Both solutions use certificates to provide authentication, message integrity, non-repudiation (through digital signatures), and privacy (encryption).
- Certificates are bound (associated) with an email address.
- To prove who has sent an email, a digital signature is added using the sender's private key. Only the sender who has the private key could have sent the message.
- To encrypt email, the message is encrypted using the recipient's public key. Only the recipient who has the private key can decrypt the message. Before you send an encrypted email to someone, you must first obtain their public key, which is normally done by having them send you a signed email.
- S/MIME uses certificates issued by either public or in-house CAs using the X.509 system.
- PGP uses two methods for validating certificates:
  - With a web of trust, individual users decide which certificates they trust. Users can then trust other designated users to introduce or recommend additional trusted users.
  - With trust signatures, digital signatures from certain certificates are trusted as being able to sign other keys. Trust signatures create a hierarchy similar to that of Certificate Authorities.
- Both S/MIME and PGP are used primarily for email encryption, although PGP can also be used for encryption of phone calls and whole disk encryption.

10.9.3 Protecting a Client from Spam (Demo Video)

Transcript:

In this demo, we're going to talk about protecting a client from spam. We're going to look at spam settings in two email services. One is Gmail, and we're going to access it through a web browser. The other is the Outlook web mail client.

Let's start with Gmail. I already have my browser open, and I've logged into Gmail. It has a few functions that are set up automatically to help you organize and sort out types of email, especially spam.
When I come down here and select this email, I get some options, including Report Spam. If I select that, it'll be moved to my Spam folder, and all emails coming from this address will be marked as spam. I'll undo that for this email.
Another thing I can do is right-click and look at some more options. A better choice for this email might be to label it Promotions rather than Spam. Now let's click on the Promotions tab and see what's in there. I see several emails, many already sorted for me. If I right-click again, I can see, under Label As, that it's labeled Promotions.
Let's look at the contents of our Spam Folder. If I click on it, we can see several emails in there. If you find some emails in here that aren't spam, you can select them and then click Not Spam. All emails in this folder are deleted 30 days after they arrive.
Now let's look at the settings. Here, you'll find Filters and Blocked Addresses. Down here, you'll find the email addresses you've previously blocked. These emails are marked as spam. If, for some reason you need to change that, this is where you do it.
Now let's look at another web-based email client, Outlook. I'm already in Outlook, and here's my inbox. I have one email down here. Let's select it. I have a tab up here called Junk. When I click on it, I get a dropdown, and I have a few choices: Junk, Phishing, and Block. Let's click each of these for more information.
First, I'll click on Junk. This will add this email to the list of Spam mail senders. I can change my mind and click Undo.
Next, click on Phishing. When we do that, this email is reported as a Phishing email to Microsoft so they can better detect these on their end. This isn't a Phishing email, so I'll click the Don't Report button.
The third option is Block. This won't report the email, but maybe it's just one that I don't care to receive any longer, and unsubscribing doesn't help. Once again, I'll cancel and not really block the email.
Let's take a look at the Junk Email folder. I have two emails in here that aren't junk, so I'll tell Outlook that, and they'll go back to my inbox. I'll do that for both emails. Now, if I look at my inbox, both emails are there.
Let's look at our settings in Outlook. I'll click on the Junk Email option here. Let's look at our options. We can block senders or entire domains from sending you email. I click on Add, and I can put in an email address or a domain that I want to block. Or maybe a sender is getting labeled as spam, but it isn't spam. If that's the case, I can add those emails and domains here, under the Safe senders and domains list. The last one here is Safe mailing lists. For example, I'm a member of several mail lists related to tech, and many spam filters mark them as spam. I can enter them here to make sure that doesn't happen.
At the bottom of the page, we have a couple of filters. The first one is to only trust emails from my Safe senders and domains list. This would be pretty restrictive and take a lot of updating. But if you only want to use this email to communicate with coworkers or family members, this might be a great thing, especially for children or the elderly. The next one will block attachments from anyone not in your Safe senders and domains list. This isn't as restrictive, and it's probably a good way to prevent malware from coming to your mailbox.
Be aware that these are just client settings, and on the server side of things, you probably have even more settings. In addition, ISPs and mail hosts are filtering emails before they even arrive to your client. In general, processes for catching spam have come a long way, but you still have to do some things manually.
That's it for this demo. In this demo, we configured spam filtering in Gmail and Outlook.

10.9.4 Securing an Email Server (Demo Video)

Transcript:

In this demonstration, we'll discuss how to prevent spam from being relayed off your email servers. To do this, we've installed an SMTP email service on this Windows Server 2022 system. Let's go to Tools.

We're going to go to the IIS 6.0 Manager. Expand that. We managed this service using the IIS Manager 6.0. Most email servers currently available will have similar SMTP capabilities and provide the same SMTP options.
Firstly, we can right-click on our SMTP server. Go to Properties. From here, we can lock down this email server. We're going to go to the Access tab. On the Access tab, you can require authentication to access this server. The default is Anonymous access, meaning anybody can connect without authenticating. You should disable this option and require at least Basic authentication.
If you have Anonymous authentication enabled and allow open relay, spammers will quickly find this system and start bouncing spam off your server to cover their tracks. If this happens, you're the one who gets in trouble, and your ISP may block your email server.
You can also use Integrated Windows Authentication, which would require you to use the domain credentials to authenticate. Basic authentication requires users just to enter a password but be aware that basic authentication is sent clear text by default, which would expose the password.
Alternatively, you can combine it with TLS encryption, the next generation of SSL, to secure authentication credentials as they're transmitted on the network. You can also require a secure connection to connect to this SMTP server in the first place.
This option requires an SSL or TLS certificate for this server to secure the connection. You can also use Connection control to configure who's allowed to connect to this SMTP server and who is blocked from connecting. This can be done using your IP addresses. You can use a range of IPs.
You can also do this using domain names as well. This requires reverse DNS lookups to be enabled. Using this information, you can restrict who is allowed access to the server. You can also implement Relay restrictions, which specify who is allowed to relay email through the SMTP server. You can allow everyone except those specified on the list or block everyone except those on the list. As before, you can also specify an IP or IP range or a domain name.
That's it for this demonstration. In this demonstration, we reviewed a few security configurations you can implement to prevent email from being relayed off of your SMTP server.

10.9.5 Configure Email Filters (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You helped your boss remove a lot of junk email, and now he would like you only to allow emails and attachments from senders on his safe sender list.

In this lab, your task is to:

Configure email filtering as follows:

- Only allow emails from the safe senders list.
- Report junk email messages to your email provider.
- Only allow attachments from the safe senders list.

Explanation

Complete this lab as follows:

Access the WebEmail filtering and reporting options.

- In the upper right corner of the WebEmail interface, select **Options** > **More Options** .
- Under *Preventing junk email* , select **Filters and reporting** .

Configure the Filters and reporting options.

Under *Choose a junk email filter* , select **Exclusive** .

Under *Report junk messages* , select **Report junk** .

Under *Block content from unknown senders* , select **Block attachments, pictures, and links for anyone not in my safe senders list** .

Select **Save** .

10.9.6 Securing Accounts on an iPad (Demo Video)

Transcript:

In this demonstration, we'll show you techniques to secure accounts on an iPad, specifically email and web browsing.

The web browser and email client on a mobile device, as with any computer, represent one of the most serious and most easily attacked vectors into the device.
Just as you need to secure the web browser and email client on a Windows workstation, you need to secure them on mobile devices as well. Let's take a look at some things you can do.
We're currently using an iPad to do this demonstration. For email, one of the key things we can do is use SSL to communicate with the mail server. This includes POP or IMAP servers, as well as the SMTP server. Let's look at how we do that.
Let's go into Settings. We'll scroll up to Passwords & Accounts Then we'll go into this Gmail account. To do that, we tap on the email address to get into the Account settings.
At the bottom, there's an Advanced selection. Tap on Advanced. You can see that we are using SSL to encrypt communications between the mail client on this mobile device and the mail server.
For this to work, the mail server for POP, IMAP, and SMTP, must be configured to support SSL. If it isn't, then turning this on won't accomplish anything.
Remember that POP3, IMAP, and SMTP transmit clear text by default. So, using SSL greatly increases the security of your email communications.
If you're using a mobile device to send email messages to someone through your mail server and you're not using SSL, then any sensitive proprietary information in that email is being transferred in clear text. It could be sniffed and read by almost anyone.
Using SSL sets up an SSL tunnel between the email client and the mail server. The SSL tunnel secures message transmission through that tunnel. It is already enabled here.
In addition to securing your email account, you also need to secure the web browser on a mobile device. Because we're using an iPad, Safari is the browser. Let's get out of the account and go to Safari Settings.
We can do several things to increase the overall security of the web browser. First, let's look at autofill. Notice that Use Contact Info is on.
To increase security, we'll turn that off. In addition, Credit Cards is also on. We'll also turn that off. This option automatically fills out credit card information. We don't want to do that; so let's keep that off.
Now we'll have to manually fill in contact and credit card information in various web forms. Autofill is a security challenge on a workstation in an office. It's a bigger problem on a mobile device.
Let's look at one more thing. We'll go back up and tap Passwords & Accounts, on the right we might want to turn off AutoFill Passwords. That way if we lose this device, someone can't access sites by having the passwords automatically populated.
We're done with this. Let's go back and go to the main Safari settings. Now we want to look at cookies. Scroll down. As you can see here under Privacy and Security, we have Cookies, and the Block All Cookies option. That's currently, that's turned off. That means that it will not block cookies. Let's turn that on.
When we turn it on, we are asked, "Hey, are you sure you want to block all cookies?" Because if we do, some websites are not going to work and the existing cookies will be removed. Let's tap Block All Cookies, and say goodbye to them.
Let's go back. Now we need to enable Fraudulent Website Warning. This is right below our Block All Cookies. This will warn us if we're visiting a site the browser thinks is a potentially fraudulent website. Right now, you can see it's turned off. Let's turn that on.

If we scroll down and tap Advanced, we can see that JavaScript is currently on. This is one of those issues where you have to balance functionality and security, just like with the cookies. Having JavaScript enabled makes the browser on this device more functional, but it also opens up security issues. The general advice is to disable JavaScript.
There will be sites that don't function properly, however, because it's commonly used in a lot of websites, but it does increase the security of the device. We can turn it off if we want to, so let's select to turn that off.
The last thing we're going to do to secure Safari is to make sure that we're blocking pop-ups. When we look here, we can see Block pop-ups is not currently on.
Let's go ahead and turn that on. That'll stop many of those annoying pop-up windows that are so common on the internet. It will help get rid of any of the annoying ads. Also, we want to prevent drive by downloads from occurring. At this point, the Safari web browser is much more secure than it was.
In this demo, we talked about things you can do to secure email and the web browser on a mobile device. In this case, an iPad. We configured the POP3 or IMAP and SMTP protocols to secure communications between the web client and the web server using SSL.
Then, we looked at what you can do to secure the Safari web browser on an iPad. We disabled autofill, we restricted cookies, we turned fraud warning on, and we turned JavaScript off. Then, we finished by turning on the pop-up blocker.

10.9.7 Secure Email on iPad (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. The receptionist, Maggie Brown, uses an iPad to manage employee schedules and messages. You need to help her secure her email and browser on her iPad.

In this lab, your task is to complete the following:

Configure Maggie's email account to use SSL for incoming mail.

Secure the internet browser as follows:

Turn off AutoFill

Turn on Block Pop-ups

Block all cookies

Turn on Fraudulent Website Warning

Turn off JavaScript

Explanation

Complete this lab as follows:

Configure email for SSL.

Select **Settings** .

Scroll down and select **Mail** .

From the right pane, select **Accounts** .

Select **Maggie Brown** .

Select **Account mbrown@gmail.com** .

Select **Advanced** .

Under *Incoming Settings* , set *Use SSL* to **ON** .

From the top, select **Account** to return to the Account menu.

Select **Done** .

Turn off AutoFill.

From the Settings menu, scroll down and select **Safari** .

From the right pane, select **AutoFill** .

Set Use Contact Info to **OFF** .

Set Names and Passwords to **OFF** .

From the top, select **Safari** to return to the Safari menu

Block all pop-ups and cookies.

From the right pane, set Block Pop-ups to **ON** .

Scroll down and set Block All Cookies to **On** .

Turn on the Fraudulent Website Warning and turn off JavaScript.

From the right pane, set Fraudulent Website Warning to **ON** .

Scroll down and select **Advanced** .

Set JavaScript to **OFF** .

10.9.8 Practice Questions (Section Quiz)

q_email_sec_blacklist_secp8

You have been receiving a lot of phishing emails sent from the domain kenyan.msn.pl. Links within these emails open new browser windows at youneedit.com.pl.

You want to make sure that these emails never reach your inbox, but you also want to make sure that emails from other senders are not affected.

What should you do?

Answers:

***Add kenyan.msn.pl to the email blacklist.**

Add pl to the email blacklist.

Add youneedit.com.pl to the email blacklist.

Add msn.pl to the email blacklist.

Explanation:

Add **kenyan.msn.pl** to the email blacklist.

Adding **msn.pl** or **pl** to the blacklist filters out all emails from kenyan.msn.pl, but this also filters out other emails from the msn.pl or pl domains.

Adding **youneedit.com.pl** to the email blacklist would prevent emails from that domain, but it would not prevent emails from kenyan.msn.pl, nor would it prevent links in the emails from opening windows to youneedit.com.pl.

q_email_sec_dkim_secp8

A cyber technician is enhancing application security capabilities for corporate email accounts following a breach.

Which of the following leverages encryption features to enable email verification by allowing the sender to sign emails using a digital signature?

Answers:

***DomainKeys Identified Mail (DKIM)**

Domain-based Message Authentication, Reporting and Conformance

Sender policy framework (SPF)

Endpoint detection and response (EDR)

Explanation:

DomainKeys Identified Mail (DKIM) leverages encryption features to enable email verification by allowing the sender to sign emails using a digital signature.

Domain-based Message Authentication, Reporting and Conformance (DMARC) checks to define rules for handling messages, such as moving messages to quarantine or spam, rejecting them outright, or tagging the message.

Sender policy framework (SPF) is an email authentication method that helps detect and prevent sender address forgery, commonly used in phishing and spam emails.

An endpoint detection and response (EDR) product's aim is not to prevent initial execution, but to provide real-time and historical visibility into the compromise, contain the malware within a single host, and facilitate remediation of the host to its original state.

q_email_sec_dmarc_01_secp8

You are the IT manager of a large corporation and have been receiving numerous complaints about phishing and spam emails.

You decide to implement a protocol that not only verifies the sender's IP address and checks the email's integrity but also defines rules for handling messages and provides reporting capabilities.

Which protocol should you implement?

Answers:

Sender Policy Framework (SPF)

DomainKeys Identified Mail (DKIM)

***Domain-based Message Authentication, Reporting and Conformance (DMARC)**

Virtual local area network (VLAN)

Explanation:

Domain-based Message Authentication, Reporting and Conformance (DMARC) is the correct answer. DMARC uses the results of SPF and DKIM checks to define rules for handling messages, such as moving messages to quarantine or spam, rejecting them outright, or tagging the message. DMARC also provides reporting capabilities, giving the owner of a domain visibility into which systems are sending emails on their behalf, including unauthorized activity. Therefore, this answer is correct.

While SPF does help in verifying the sender's IP address against a list of authorized sending IP addresses, it does not provide the capability to define rules for handling messages or provide reporting capabilities. Therefore, this answer is incorrect.

DKIM allows the sender to sign emails using a digital signature, which the receiving email server can verify. However, like SPF, it does not provide the capability to define rules for handling messages or provide reporting capabilities. Therefore, this answer is incorrect.

A VLAN is a network topology configuration for creating distinct broadcast domains. It is not a protocol for email security or spam prevention. Therefore, this answer is incorrect.

q_email_sec_dmarc_02_secp8

A technician is modifying controls to increase securities on messaging services.

Which of the following options check to define rules for handling messages, such as moving messages to quarantine or spam, rejecting them outright, or tagging the message?

Answers:

***Domain-based Message Authentication, Reporting and Conformance (DMARC)**

Sender policy framework (SPF)

Endpoint detection and response (EDR)

DomainKeys Identified Mail (DKIM)

Explanation:

Domain-based Message Authentication, Reporting and Conformance (DMARC) checks to define rules for handling messages, such as moving messages to quarantine or spam, rejecting them outright, or tagging the message.

Sender policy framework (SPF) is an email authentication method that helps detect and prevent sender address forgery, commonly used in phishing and spam emails.

An endpoint detection and response (EDR) product's aim is not to prevent initial execution, but to provide real-time and historical visibility into the compromise, contain the malware within a single host, and facilitate remediation of the host to its original state.

DomainKeys Identified Mail (DKIM) leverages encryption features to enable email verification by allowing the sender to sign emails using a digital signature.

q_email_sec_email_secp8

You install a new Linux distribution on a server in your network. The distribution includes a Simple Mail Transfer Protocol (SMTP) daemon that is enabled by default when the system boots. The SMTP daemon does not require authentication to send email messages.

Which type of email attack is this server susceptible to?

Answers:

***Open SMTP relay**

Phishing

Sniffing

Viruses

Explanation:

An SMTP relay is an email server that accepts mail and forwards it to other mail servers, and an open SMTP relay allows anyone to forward mail if they choose. If your mail server is an open SMTP relay, spammers can also take advantage of it to obscure the actual source of the email. If spammers use your relay for sending mail, your server may soon be placed on a blacklist. Other mail servers will then stop receiving any mail (even legitimate mail) sent from your servers. As a best practice:

Configure your mail server to accept mail only from authenticated users or specific email servers that you authorize.

Require TLS encryption to connect to the server.

A phishing scam uses an email pretending to be from a trusted organization that asks you to verify personal information or send money.

Sniffing occurs when a user captures packets from a network and inspects their contents.

Viruses are types of malware that spread by infecting legitimate files on a computer system and are sometimes sent as email attachments.

q_email_sec_pgp_secp8

You are the IT security manager for a global company that frequently exchanges sensitive information via email.

The company has been using Secure/Multipurpose Internet Mail Extensions (S/MIME) for email security, but due to recent cyber threats, you are considering implementing an additional layer of security.

Which of the following options would provide the MOST robust solution for securing your company's email communications?

Answers:

Implementing DomainKeys Identified Mail (DKIM)

Implementing Sender Policy Framework (SPF)

***Implementing Pretty Good Privacy (PGP)**

Implementing Domain-based Message Authentication, Reporting and Conformance (DMARC)

Explanation:

Implementing Pretty Good Privacy (PGP) is the correct answer. PGP provides authentication, message integrity, non-repudiation (through digital signatures), and privacy (encryption). It uses a combination of symmetric and asymmetric encryption to secure email communication, making it the most robust solution among the options provided.

Implementing DomainKeys Identified Mail (DKIM) is an email authentication method that helps detect email spoofing by allowing the receiver to check if the email was sent and signed by the domain it claims to be from. However, it does not provide encryption for the content of the email, making it less secure than PGP.

Implementing Sender Policy Framework (SPF) is an email authentication method that helps detect and prevent sender address forgery. While SPF can help prevent phishing and spam emails, it does not provide encryption for the content of the email, making it less secure than PGP.

Implementing Domain-based Message Authentication, Reporting and Conformance (DMARC) uses the results of SPF and DKIM checks to define rules for handling messages. While DMARC can help prevent phishing and spam emails, it does not provide encryption for the content of the email, making it less secure than PGP.

q_email_sec_phishing_01_secp8

Users in your organization receive email messages informing them that suspicious activity has been detected on their bank accounts. They are directed to click a link in the email to verify their online banking username and password. The URL in the link is in the .ru top-level DNS domain.

Which kind of attack has occurred?

Answers:

***Phishing**

Open SMTP relay

Virus

Buffer overflow

Explanation:

A phishing scam uses an email that purports to be from a trusted organization and asks you to verify personal information or send money. In a phishing attack:

- A fraudulent message (which appears to be legitimate) is sent to a target.
- The message requests that the target visit a fraudulent website (which also appears to be legitimate). Graphics, links, and web pages look almost identical to legitimate requests from legitimate websites.
- The fraudulent website requests that the victim provide sensitive information, such as an account number and password.

An SMTP relay is an email server that accepts mail and forwards it to other mail servers.

In a buffer overflow attack, a program (while writing data to a memory buffer) overruns the buffer's boundaries and writes data in adjacent memory addresses.

Viruses are types of malware that spread by infecting legitimate files on a computer system and are sometimes sent as email attachments.

q_email_sec_phishing_02_secp8

Which of the following BEST describes phishing?

Answers:

- ***A fraudulent email that claims to be from a trusted organization.**
- Malware that often uses email as its distribution mechanism.
- Unwanted and unsolicited email sent to many recipients.
- An email server that accepts mail and forwards it to other mail servers.

Explanation:

Phishing is a fraudulent email that claims to be from a trusted organization.

Spam is unwanted and unsolicited email sent to many recipients.

A virus is malware that often uses email as its distribution mechanism.

An SMTP relay is an email server that accepts mail and forwards it to other mail servers.

q_email_sec_phishing_03_secp8

Which of the following would you do to help protect against phishing?

Answers:

- ***Only open emails if you recognize the sender.**
- Don't post your full email address anywhere on the web.

Don't click on an unsubscribe link at the bottom of an unsolicited email.

In the email client, disable preview screens.

Explanation:

To protect against phishing:

Check the email header information to see more info about the sender and the links that are in the email.

Only open emails if you recognize the sender.

Check the actual link destination within emails to verify that they go to the correct URL and not a spoofed one.

Do not click on links in emails. Instead, type the real URL into the browser. You could also look up the website in a search engine.

Verify that HTTPS is used when going to e-commerce sites. HTTPS requires a certificate that matches the server name in the URL that is verified by a trusted certificate authority (CA). You can also look for the lock icon to verify that HTTPS is used.

Implement phishing protections within your browser.

To control spam:

Enable spam filters on client and email servers. Filter junk email by identifying safe senders (whitelists), blocked senders (blacklists), countries to block email from, and languages to block.

Enable antivirus scanning for attachments on the client and email servers.

In the email client, disable preview screens. An email can have links for active items that can report back to the spammer.

Don't click on an unsubscribe link at the bottom of an unsolicited email. Doing this verifies to the spammer that the email address is a current and active email address. Only unsubscribe from trusted organizations.

Install server-level, anti-spam software on the email server.

Don't post your full email address anywhere on the web. Spammers use software to scan websites to find email addresses and then add them to their email lists for spamming.

q_email_sec_phishing_04_secp8

An accounts payable clerk received an email requesting payment information for materials for an ongoing project. The email appears to be from a known vendor.

Before giving the information over, what should the clerk protect against?

Answers:

Typosquatting

Coercion

***Phishing**

Consensus technique

Explanation:

Phishing is a combination of social engineering and spoofing. It persuades or tricks the target into interacting with a malicious resource disguised as a trusted one.

Typosquatting means that the threat actor registers a domain name very similar to a real one, hoping that users will not notice the difference and assume they are browsing a trusted site.

Coercion or the use of urgency refers to the intimidation of the target with a bogus appeal to authority or penalty, such as getting fired or not acting quickly enough to prevent some dire outcome.

The persuasive or consensus technique attempts to convince the target that the request is natural and it would be impolite or somehow odd to refuse.

q_email_sec_smime_secp8

Which of the following mechanisms can you use to add encryption to email? (Select two.)

Answers:

***PGP**

***S/MIME**

HTTPS

Reverse DNS

Secure Shell

Explanation:

Use Pretty Good Privacy (PGP) or Secure MIME (S/MIME) to add encryption to emails.

HTTPS is used by web browsers to request data from web servers.

Secure Shell (SSH) is a secure remote management utility.

Reverse DNS can be used to verify the sending device's IP address included in an email. However, this does not add encryption to email messages.

q_email_sec_spam_01_secp8

If an SMTP server is not properly and securely configured, it can be hijacked and used maliciously as an SMTP relay agent.

Which activity could result if this happens?

Answers:

Salami attack

Data diddling

Virus hoax

***Spamming**

Explanation:

Attackers often distribute spam by hijacking a misconfigured SMTP server. SMTP servers that act as relay agents for unauthorized or external users can be easily employed to deliver spam. It is extremely important to properly configure SMTP servers to accept email only from authorized internal users.

A salami attack is an attack where small amounts of information, data, or valuables are taken over a period of time. The result is to construct or obtain data or property of great value. A common example of a salami attack is to deposit the fractions of cents from an accounting program into a numbered account. Eventually, the fraction deposits total a significant sum.

Data diddling is changing information during input, processing, output, or storage.

A virus hoax is a social engineering attack designed to play off of the fears of victims to convince them to perform malicious activities against themselves.

q_email_sec_spam_02_secp8

Which type of malicious activity can be described as numerous unwanted and unsolicited email messages sent to a wide range of victims?

Answers:

Brute force

Trojan horse

Hijacking

***Spamming**

Explanation:

Spamming is a type of malicious activity can be described as numerous unwanted and unsolicited email messages being sent to a wide range of victims. Spam itself is not usually malicious in nature. More often than not, it is advertising for some product or service. Unfortunately, spam accounts for 40 to 60 percent of all email traffic on the internet. Most of this activity is unsolicited.

A brute force attack is a hacking method that uses trial and error to crack passwords, login credentials, and encryption keys. It does not involve sending numerous unwanted and unsolicited email messages sent to a wide range of victims.

A trojan horse is a type of malware that downloads onto a computer disguised as a legitimate program. It does not involve sending numerous unwanted and unsolicited email messages sent to a wide range of victims.

Hijacking (or computer hijacking) is a type of network security attack in which the attacker takes control of computer systems, software programs and/or network communications. It does not involve sending numerous unwanted and unsolicited email messages sent to a wide range of victims.

q_email_sec_spam_03_secp8

An attacker sends an unwanted and unsolicited email message to multiple recipients with an attachment that contains malware.

Which kind of attack has occurred in this scenario?

Answers:

***Spam**

Open SMTP relay

Phishing

Repudiation attack

Explanation:

Spam is unwanted and unsolicited email messages sent to many recipients. Spam:

Can be benign, such as emails trying to sell products.

Can be malicious, such as emails containing phishing content, drive-by downloads, or malware.

Can contain malware as attachments.

Wastes bandwidth and could fill an inbox, resulting in a denial-of-service condition.

An open SMTP relay allows anyone to forward mail. An open SMTP relay can be used by spammers to send mail.

A phishing scam is an email pretending to be from a trusted organization, asking the recipient to verify personal information or send money.

In a repudiation attack, an attacker accesses your email server and sends spoofed emails to others, making them appear as if they came from you.

q_email_sec_spf_secp8

An organization is increasing security on corporate email exchanges after being a target in a whaling campaign.

Which of the following options is an email authentication method that helps detect and prevent sender address forgery?

Answers:

***Sender Policy Framework (SPF)**

Endpoint detection and response (EDR)

DomainKeys Identified Mail (DKIM)

Domain-based Message Authentication, Reporting and Conformance (DMARC)

Explanation:

Sender Policy Framework (SPF) is an email authentication method that helps detect and prevent sender address forgery, commonly used in phishing and spam emails.

An endpoint detection and response (EDR) product's aim is not to prevent initial execution, but to provide real-time and historical visibility into the compromise, contain the malware within a single host, and facilitate remediation of the host to its original state.

DomainKeys Identified Mail (DKIM) leverages encryption features to enable email verification by allowing the sender to sign emails using a digital signature.

Domain-based Message Authentication, Reporting and Conformance (DMARC) checks to define rules for handling messages, such as moving messages to quarantine or spam, rejecting them outright, or tagging the message.

11.0 Security Governance Concepts

11.1 Policies, Standards, and Procedures

As you study this section, answer the following questions:

What role do policies play in the framework of an organization?

Why is compliance to policies important?

How are standards different from policies?

Describe four different types of standards an organization may establish internally.

What problems does a playbook help an organization address?

The key terms for this section include:

Term	Definition
Policies	A strictly enforceable rule set that determines how a task should be completed.
Acceptable Use Policy (AUP)	A policy that governs employees' use of company equipment and Internet services. ISPs may also apply AUPs to their customers.
Information Security Policies	A document or series of documents that are backed by senior management and that detail requirements for protecting technology and information assets from threats and misuse.
Business Continuity & Continuity of Operations Plans (COOP)	A collection of processes that enable an organization to maintain normal business operations in the face of some adverse event.
Disaster Recovery	A documented and resourced plan showing actions and responsibilities to be used in response to critical incidents.
Software Development Life Cycle (SDLC)	The processes of planning, analysis, design, implementation, and maintenance that often govern software and systems development.
Guidelines	Best practice recommendations and advice for configuration items where detailed, strictly enforceable policies and standards are impractical.

Standards	Expected outcome or state of a task that has been performed in accordance with policies and procedures. Standards can be determined internally or measured against external frameworks.
Procedures	Detailed instructions for completing a task in a way that complies with policies and standards.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.6 Given a scenario, implement and maintain identity and access management. Multifactor authentication Password concepts Length Complexity Reuse Expiration 5.1 Summarize elements of effective security governance. Guidelines Policies Acceptable use policy (AUP) Information security policies Business continuity Disaster recovery Incident response Software development lifecycle (SDLC) Change management Standards

	Password Access control Physical security Encryption Procedures Playbooks External considerations Regulatory Legal Industry Local/regional National Global Monitoring and revision Types of governance structures Boards Committees Government entities Centralized/decentralized

11.1.1 Program Management and Oversight Overview (Lesson Video)

Transcript:

A critical component of any Information Security plan is proper program management and oversight. This ensures that all security controls put in place are being implemented and working properly. In this lesson, we'll look at some of the key elements of effective security governance, including organizational standards, external considerations, and ongoing monitoring and revision of the security plan.

A proper security plan will include policies, procedures, guidelines, and standards. Policies are essentially organizational rules that must be followed by employees. A common example is an Acceptable Use Policy. Procedures outline how

certain functions should be performed. Guidelines are designed to provide support and additional information on policies and procedures. Standards define how the policies are implemented.
For example, an organization may have a policy that access control must be implemented on certain data types. The organization would then implement standards that define who can access the data, how they can access it, and how access is monitored.
Password policies are typically defined in an organization. The standards define the length and complexity of the password, how often a password can be reused, and how often the password should be changed.
For example, the standards might state that the password should be a minimum of 8 characters, including upper and lowercase, numbers, and symbols, and that there should not be a set time to change the password. The standard can even define that multi-factor authentication should be used.
Some other common standards an organization should have defined include physical security measures such as cameras, physical access control, and security guard schedules and placements. To protect sensitive data, encryption standards should also be defined and implemented throughout the organization.
External considerations must be looked at when establishing internal policies. Perhaps the biggest external consideration is any laws that apply to the organization.
For example, in the health industry, the Health Insurance Portability and Accountability Act, or HIPAA, establishes laws and regulations that must be followed by any health organization. These laws and regulations define how data is secured, who can access patient data, and more. When internal policies are established, they must adhere to the regulations established by the law.
Aside from regulatory and legal compliance, we also need to consider any industry standards that may be in place. For example, in the IT field, the National Institute of Standards and Technology, or NIST, publishes various industry standards. While these aren't laws or regulations, these industry standards are put into place to define specific practices that should promote performance and efficiency. While some standards are globally accepted, others differ between countries, states, and even between cities.
Monitoring and updating of all policies and standards should take place on a regular basis. As new technologies come out, new laws and industry standards will be released. It's up to the organization to review these and update their internal policies and standards as needed.
That'll wrap up this lesson on program management and oversight. In this lesson, we first went over the differences in standards, policies, procedures, and guidelines. Standards define how policies should be implemented and are a core component of any security governance plan. We then looked at some external factors that should be considered when developing a security plan. These include legal and regulatory compliance, industry standards, and even global, national, and local standards and practices. We wrapped up by going over why it's important to monitor and revise the plans to stay current with industry and regulatory changes.

11.1.2 Program Management and Oversight Facts

Security governance is a critical aspect of an organization's overall security posture, providing a framework that guides the management of cybersecurity risks. It involves developing, implementing, and maintaining policies, procedures, standards, and guidelines to safeguard information assets and technical infrastructure. Security governance encompasses the roles and responsibilities of various stakeholders, emphasizing the need for a culture of security awareness throughout the organization. Governance frameworks must manage and maintain compliance with relevant laws, regulations, and contractual obligations while supporting the organization's strategic objectives. Effective security governance also involves continuous monitoring and improvement to adapt to evolving threats and changes in the business and regulatory environment.

This lesson covers the following topics:

- Governance and accountability
- Governance boards
- Government entities and groups
- Legal Environment

Governance and Accountability

Governance practices ensure organizations abide by all applicable cybersecurity laws and regulations to protect them from legal liability. Governance and organizational-level oversight must manage many legal risks, such as regulatory compliance requirements, contractual obligations, public disclosure laws, breach liability, privacy laws, intellectual property protection, and licensing agreements, and interpret and translate these legal requirements into operational controls to avoid legal trouble, act ethically, and protect the organization.

Guidelines for Implementing Governance

Follow these guidelines to support effective governance controls:

- Implement a governance structure that best supports the organization's objectives.
- Leverage expertise through committees to support decision-making.
- Establish a comprehensive list of policies, processes, standards, and guidelines.
- Implement change management programs to maintain control and promote transparency.
- Use automation and orchestration tools to improve consistency, reduce response times, and support compliance.

Governance Boards

Governance boards are crucial in ensuring an organization's effective security governance and oversight because they are responsible for setting strategic objectives, policies, and guidelines for security practices and risk management. Governance boards oversee the implementation of security controls, work closely with risk management teams to ensure compliance with relevant laws and regulations, and evaluate the security program's overall effectiveness. Governance boards drive organization-wide security practices through leadership, guidance, and accountability and ensure that security risks are effectively identified and mitigated. Governance boards unite executive management, security professionals, and stakeholders to ensure security is a top strategic priority aligned with the organization's objectives and values.

Centralized versus Decentralized

Centralized and decentralized security governance models aim to achieve the organization's security goals, protect assets, mitigate risks, and ensure regulatory compliance. Additionally, they recognize the importance of security and the need for collaboration between stakeholders and departments. However, there are notable differences between the two approaches. In centralized security governance, decision-making authority primarily rests with a single core group or department that establishes policies, procedures, and guidelines and makes important security-focused decisions. Resource allocation, including budget and personnel, is controlled by this group to promote consistency and standardization across the entire organization.

In contrast, decentralized security governance distributes decision-making authority to different groups or departments to facilitate security-focused decisions based on localized needs and priorities. Each unit has greater control over the allocation of security resources to allow greater adaptability and tailoring of security capabilities.

The choice between centralized and decentralized security governance depends on the organization's size, structure, culture, and risk appetite. Ultimately, the goal is to create a security governance model that effectively supports the organization's needs while balancing security risks.

Hybrid governance structures combine elements of both centralized and decentralized approaches. It aims to balance the advantages of centralized oversight and decentralized implementation. Under a hybrid system, specific security processes and

decisions are centralized, while others are delegated to business units or departments to facilitate the development of standardized policies at the enterprise level while providing flexibility and local control as warranted.

Committees and Boards

Governance boards depend upon governance committees to assist in complex decision-making situations. The governance board is typically composed of executives with the ultimate decision-making authority and is responsible for setting the strategic direction and policies of the organization. This responsibility often requires executives to make critical decisions regarding subjects outside their scope of expertise.

Committees are specialized groups comprised of subject matter experts, stakeholders, and representatives from relevant departments that focus on specific issues, such as security, risk management, audit, or compliance. They provide in-depth analysis, recommendations, and operational support to the governance board to provide them with the critical information needed to make effective decisions.

Governance boards and governance committees serve distinct roles within an organization's governance structure. Governance boards are typically composed of high-level executives and external stakeholders; whereas, governance committees are typically comprised of subject matter experts and operational leaders.

Data Governance Roles

Security governance relies heavily on specially designed and interdependent roles: owner, controller, processor, and custodian. Each role carries unique responsibilities that contribute to maintaining effective security oversight and control.

Role	Description
Owner	A high-ranking employee, like a director or a vice president, typically holds the owner role and is ultimately responsible for ensuring data is appropriately protected. The owner identifies what level of classification and sensitivity the data has, decides who should have access to it, and what level of security should be applied. In relation to governance, the owner provides strategic guidance to ensure that security policies align with business objectives.
Controller	The controller role closely relates to General Data Protection Regulation (GDPR) and identifies the purposes, conditions, and means of processing personal data. An individual, public authority, agency, or other body can fill the controller role. The controller ensures that data processing activities adhere to all legal requirements. In relation to governance, the controller helps maintain legal and regulatory compliance.
Processor	The processor is responsible for processing personal data on behalf of the controller and often represents cloud service providers (CSP) but could also be represented by vendors and business partners. Processors must maintain records of their processing activities, cooperate with supervisory authorities, and implement appropriate security measures to protect the data they handle. In relation to governance, the processor role ensures that data is handled securely and in accordance with the rules established by the owner and controller roles.
Custodian	The custodian, also known as the data steward, is responsible for the safe custody, transport, storage of the data, and implementation of business rules. The IT department typically represents the custodian role, and in relation to governance, the custodian role implements and enforces the security controls established by the data owner and controller and reports any issues indicative of a security incident.

Coordination among data owner, controller, processor, and custodian in managing and protecting data is crucial to ensure compliance with data protection regulations, establish clear responsibilities, and maintain data integrity and security.

Government Entities and Groups

At the government level, governance committees are often represented by specialized agencies. Several government agencies are associated with security governance and differ between countries and jurisdictions. A few examples of government agencies with security governance responsibilities include the following:

Agency	Description
Regulatory agencies	Regulatory agencies establish and enforce security standards, regulations, and guidelines. They oversee compliance with laws related to specific sectors such as finance, healthcare, telecommunications, and energy.
Intelligence agencies	Intelligence agencies gather and analyze information to identify and counteract potential security threats and provide this information to national-level government groups to steer national policy and military strategy.
Law enforcement agencies	Law enforcement agencies enforce laws and regulations related to public safety and security. They investigate and prosecute criminal activities, including cybercrimes and terrorist activities.
Defense and military organizations	Defense and military organizations are responsible for safeguarding national security and protecting the country from external threats. They develop strategies, policies, and capabilities to address physical security, border control, and defense-related cybersecurity.
Data protection authorities	Data protection authorities focus on protecting personal data and privacy rights. They enforce data protection regulations and provide guidance on the best practices for securing personal information.
National cybersecurity agencies	National cybersecurity agencies focus on protecting critical infrastructure, government networks, and national cybersecurity interests. They develop cybersecurity strategies, coordinate incident response, and provide guidance on cybersecurity practices for government entities and private organizations.

Legal Environment

Governance committees ensure their organizations abide by all applicable cybersecurity laws and regulations to protect them from legal liability. The governance committee must address these external considerations in the strategic plan for the organization.

Governance committees must manage many legal risks, such as regulatory compliance requirements, contractual obligations, public disclosure laws, breach liability, privacy laws, intellectual property protection, licensing agreements, and many others. Cybersecurity governance committees must interpret and translate these legal requirements into operational controls to avoid legal trouble, act ethically, and protect the organization.

The key frameworks, benchmarks, and configuration guides may be used to demonstrate compliance with a country's legal/regulatory requirements or with industry-specific regulations. *Due diligence* is a legal term meaning that responsible persons have not been negligent in discharging their duties. Negligence may create criminal and civil liabilities. Many countries have enacted legislation that criminalizes negligence in information management. In the United States, for example, the Sarbanes-Oxley Act (SOX) mandates the implementation of risk assessments, internal controls, and audit procedures. The Computer Security Act (1987) requires federal agencies to develop security policies for computer systems that process confidential

information. In 2002, the Federal Information Security Management Act (FISMA) was introduced to govern the security of data processed by federal government agencies.

Some regulations have specific cybersecurity control requirements; others simply mandate "best practice," as represented by a particular industry or international framework. It may be necessary to perform mapping between different industry frameworks, such as NIST and ISO 27K, if a regulator specifies the use of one but not another. Conversely, the use of frameworks may not be mandated as such, but auditors are likely to expect them to be in place as a demonstration of a strong and competent security program.

Global Law

As information systems become more interconnected globally, many countries have enacted laws with broader, international reach. Some examples include the General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA) which both act to protect the privacy of the constituents associated with each respective law irrespective of geopolitical boundaries.

Personal Data and the General Data Protection Regulation (GDPR)

Where some types of legislation address cybersecurity due diligence, others focus in whole or in part on information security as it affects privacy or personal data. Privacy is a distinct concept and requires that collection and processing of personal information be both secure and fair. Fairness and the right to privacy, as enacted by regulations such as the European Union's General Data Protection Regulation (GDPR) , means that personal data cannot be collected, processed, or retained without the individual's informed consent. *Informed consent* means that the data must be collected and processed only for the stated purpose, and that purpose must be clearly described to the user in plain language, not legal jargon. GDPR (ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr) gives data subjects rights to withdraw consent, and to inspect, amend, or erase data held about them. Failure to comply with GDPR rules can result in incredibly large fines.

California Consumer Privacy Act (CCPA)

The CCPA provides California residents the right to know what personal information businesses collect about them, the purpose of collecting this data, and with whom they share it. It protects California residents' rights to access their personal information, delete it, or opt out of its sale. Organizations must inform consumers about the categories of personal information they collect and the purposes for which the information will be used. The CCPA applies to any organization, regardless of its location, that provides goods or services to California residents; has gross annual revenues over $25 million; buys or sells the personal information of 50,000 or more consumers, households, or devices; or derives 50% or more of annual revenues from selling personal information.

Varonis's blog contains a useful overview of privacy laws in the United States (varonis.com/blog/us-privacy-laws).

Regulations and National, Local, Regional and Industry Laws

Many countries have national-level laws to support effective cybersecurity practices and protect citizen data. The scope and detail of these laws varies significantly from one country to another, but organizations must comply with the laws in all the jurisdictions where they operate.

Examples in the United States include the Health Insurance Portability and Accountability Act (HIPAA), Gramm-Leach-Bliley Act (GLBA), and Federal Information Security Management Act (FISMA). The United Kingdom's national laws include the Data Protection Act 2018 and the Network and Information Systems (NIS) Regulations 2018. Canada enforces a privacy law called Personal Information Protection and Electronic Documents Act (PIPEDA), India regulates the IT industry with the Information Technology Act 2000, and Australia enforces the Privacy Act 1988. While it is not necessary to understand the details of each of these laws, it is important to understand that this complicated legal framework greatly influences the functional elements of a cybersecurity program.

While most cybersecurity laws typically have national or international scope due to the global nature of the internet, there are also laws and regulations with a more local or regional reach. These laws may be specific to states, provinces, or even cities, particularly in larger countries like the United States. Two important examples include the New York Department of Financial Services (DFS) Part 500 Cybersecurity Regulation and the Massachusetts 201 CMR 17.00 .

Industry-specific cybersecurity laws and regulations govern how data should be handled and protected. Here are a few key examples to help highlight that cybersecurity is a significant concern across all sectors of industry and is protected by a complicated matrix of laws that must be accommodated by cybersecurity operations and organizational governance:

Description of the Table

Sector	Laws
Healthcare	Health Insurance Portability and Accountability Act (HIPAA) (United States) The General Data Protection Regulation (GDPR) (European Union)
Financial Services	Gramm-Leach-Bliley Act (GLBA) (United States) Payment Card Industry Data Security Standard (PCI DSS) (Contractual obligation)
Telecommunications	Communications Assistance for Law Enforcement Act (CALEA) (United States)
Energy	North American Electric Reliability Corporation (NERC) (United States and Canada)
Education & Children	Family Educational Rights and Privacy Act (FERPA) (United States) Children's Internet Protection Act (CIPA) (United States) Children's Online Privacy Protection Act (COPPA) (United States)
Government	Federal Information Security Modernization Act (FISMA) (United States) Criminal Justice Information Services Security Policy (CJIS) (United States) The Government Security Classifications (GSC) (United Kingdom)

Cybersecurity regulations are legal rules and guidelines formulated by governments and regulatory bodies to safeguard digital information and systems from cyber threats. They set standards for protecting data confidentiality, integrity, and availability, particularly sensitive and personal information. Regulations cover diverse capabilities, including data protection, network and information systems security, data breach notifications, digital identity verification, and many others. Businesses, government agencies, organizations, and executives must work diligently to comply with these regulations or risk significant fines and imprisonment.

Cybersecurity regulations aim to protect consumer privacy rights, ensure the security of the financial system, uphold the stability and trustworthiness of the Internet and digital economy, and protect critical national infrastructure from cybercrime. The applicability of regulations depends on factors such as the industry the organization operates within, the types of data it handles, and the regions where it conducts business. Here are a few key examples, several of which have been previously mentioned:

General Data Protection Regulation (GDPR)

California Consumer Privacy Act (CCPA)

Health Insurance Portability and Accountability Act (HIPAA)

Federal Information Security Management Act (FISMA)

Network and Information Systems (NIS) Directive

Cybersecurity Maturity Model Certification (CMMC)

Monitoring and Revision

The cybersecurity landscape is continually evolving. Consequently, organizations must ensure that their cybersecurity policies, procedures, standards, and legal compliance practices associated with legal and regulatory compliance are regularly monitored, evaluated, and updated. These responsibilities are generally managed via collaboration among diverse groups to review existing policies, procedures, and standards and ensure their effectiveness against current requirements. Routine audits, inspections, and assessments are commonly used to measure compliance levels and identify new risks. The results of compliance reports, technological changes, business processes, laws, or newly identified risks drive policy, procedure, and standards revisions. Regular training sessions help to inform employees of policy changes and ensure continued compliance.

Additionally, organizations must maintain awareness of any changes in cybersecurity legislation in their jurisdictions, including international, national, regional, or industry-specific laws. Effective monitoring and revision of cybersecurity policies, procedures, standards, and legal compliance practices is a dynamic, cyclical process requiring diligence, foresight, and proactive strategies.

11.1.3 Policies, Standards, and Procedures

Policies, standards, and procedures are three key components that form the foundation of an organization's security program. Policies are high-level, authoritative documents defining the organization's security commitment. Standards are more specific than policies and specify the methods used to implement technical and procedural requirements. Procedures are detailed, step-by-step instructions describing how to complete specific tasks and align to the requirements provided in standards. Procedures provide clear directions for individuals to perform their job duties consistently, securely, and efficiently.

This lesson covers the following topics:

Policies

Standards

Procedures

Policies

Organizational policies are vital in establishing effective governance and ensuring organizational compliance. They form the framework for operations, decision-making, and behaviors, setting the rules for a compliant and ethical corporate culture. Governance describes the processes used to direct and control an organization, including the processes for decision-making and risk management. Policies are the outputs of governance. They establish the rules that frame decision-making processes, risk mitigation, fairness, and transparency. They set expectations for performance, align the organization around common goals, prevent misconduct, and remove inefficiencies.

Compliance describes how well an organization adheres to regulations, policies, standards, and laws relevant to its operation. Organizational policies are critical in ensuring compliance by integrating legal and regulatory requirements into daily operations. Policies define the rules and procedures for maintaining compliance and outline the consequences of noncompliance.

For example, an organization may have a data privacy policy that explains how it will maintain compliance with relevant laws to protect customer data. The policy details data collection, storage, processing, and sharing practices, including employee responsibilities, to ensure that all organization members understand and adhere to the rules. Organizational policies help facilitate compliance assessments through internal and external audits as policies provide a roadmap auditors follow to determine whether an organization is operating as it claims and is successfully satisfying its regulatory obligations.

Common Organizational Policies

Acceptable Use Policy (AUP) — This policy outlines the acceptable ways in which network and computer systems may be used by defining what constitutes acceptable behavior by users. AUPs typically address browsing behavior, appropriate content, software downloads, and handling sensitive information. The goal of an AUP is to ensure that users do not engage in activities that could harm the organization or its resources. Also, the AUP should detail the consequences for non compliance, including details regarding how compliance is monitored and require employees to acknowledge their comprehension of the AUP's rules via signature.

Information Security Policies — These are policies created by an organization to ensure that all information technology users comply with rules and guidelines related to the security of the information stored within the environment or the organization's sphere of authority.

Business Continuity & Continuity of Operations Plans (COOP) — Business continuity and COOP policies focus on the critical processes that must remain operational during and after a substantial disruption, like a natural disaster or a cyberattack.

Disaster Recovery — These policies detail the steps required to recover from a catastrophic event such as a natural disaster, major hardware failure, or a significant security breach. The goal is to restore operations as quickly and efficiently as possible.

Incident Response —This policy outlines the processes to be followed after a security breach or cyberattack occurs. It details the steps for identifying, investigating, controlling, and mitigating the impact of incidents, including procedures for communicating about the incident to internal and external sources.

Software Development Life Cycle (SDLC) —SDLC policies govern software development within an organization. These policies provide a structured plan detailing the stages of development from initial requirement analysis to maintenance after deployment. It ensures that all software produced meets the organization's efficiency, reliability, and security standards.

Change Management —Change management policies outline how changes to IT systems and software are requested, reviewed, approved, and implemented, including all documentation requirements.

Guidelines

Guidelines describe recommendations that steer actions in a particular job role or department. They are more flexible than policies and allow greater discretion for the individuals implementing them. Guidelines provide best practices and suggestions on achieving goals and completing tasks effectively and help individuals understand the required steps to comply with a policy or improve effectiveness.

An example of a guideline might be related to help desk support practices related to using email in response to employee support requests. The guideline may recommend specific language, tone, or response times but would allow for flexibility depending on the request's circumstances. While both policies and guidelines work to steer the actions and behaviors of employees, policies are mandatory and define strict rules, whereas guidelines provide recommendations and allow for more individual judgment and discretion. Regular review of guidelines is important to ensure they remain practical and relevant. Periodic assessments and updates to guidelines allow organizations to adapt them to changing technologies, business operations, emerging threats, and evolving industry standards.

Standards

Standards define the expected outcome of a task, such as a particular configuration state for a server, or performance baseline for a service. The selection and application of standards within an organization center on various dynamic elements such as regulatory requirements, business-specific needs, risk management strategies, industry practices, and stakeholder expectations.

Regulatory requirements are the primary driver for adopting standards. The unique operational differences between organizations dictate varying legal requirements and security, privacy, and data protection regulations. These requirements often require implementing specific standards or using guidelines for achieving compliance. The healthcare industry in the United States is a classic example, where providers must comply with stringent data protection and privacy standards established by the Health Insurance Portability and Accountability Act (HIPAA).

Depending on the nature of its operations, customer base, or technological dependencies, each organization must adopt standards that specifically address its needs. For example, organizations heavily utilizing credit card transactions will adopt the PCI DSS standard to safeguard the cardholder data environment (CDE). Similarly, cloud-reliant organizations often prefer adopting ISO/IEC 27017 and ISO/IEC 27018 to ensure safe and secure cloud operations.

Risk management strategies organizations stress the need for appropriate standards. Standards help identify, evaluate, and manage risks and fortify the organization's resilience against security incidents or data breaches. ISO/IEC 27001, for example, provides a comprehensive framework for an information security management system (ISMS) designed to aid organizations in effectively managing security risks. Adherence to industry best practices also influences the adoption of standards. Conforming to widely accepted and tested standards demonstrates an organization's commitment to upholding high security and data protection levels to bolster the organization's reputation and build trust with customers and partners. Stakeholder expectations (such as customers, partners, vendors, investors, executive boards, etc.) significantly influence the choice of standards too. Stakeholders view adherence to recognized standards as an affirmation of the organization's dedication to quality, security, and reliability.

The choice of standards should not be a procedural decision but instead a strategic one. The selection of standards involves a thoughtful balance of legal and regulatory requirements, business-specific needs, risk management protocols, industry best practices, and stakeholder expectations. Adopting standards impacts how a business operates, and selecting appropriate standards helps an organization run more effectively. In contrast, adopting the wrong standards, or failing to plan the implementation of standards properly, can have severe negative consequences.

Industry Standards

Common industry standards used by public and private organizations include the following:

ISO/IEC 27001 —An international standard that provides an i nformation s ecurity m anagement s ystem (ISMS) framework to ensure adequate and proportionate security controls are in place.

ISO/IEC 27002 —This is a companion standard to ISO 27001 and provides detailed guidance on specific controls to include in an ISMS.

ISO/IEC 27017 —An extension to ISO 27001 and specific to cloud services.

ISO/IEC 27018 —Another addition to ISO 27001, and specific to protecting personally identifiable information (PII) in public clouds.

NIST (National Institute of Standards and Technology) Special Publication 800-63 —A U S government standard for digital identity guidelines, including password and access control requirements.

PCI DSS (Payment Card Industry Data Security Standard) —A standard for organizations that handle credit cards from major card providers, including requirements for protecting cardholder data.

FIPS (Federal Information Processing Standards) —FIPS are standards and guidelines developed by NIST for federal computer systems in the U nited States that specify requirements for cryptography.

Common industry standards such as these play a significant role in auditing by providing a benchmark for evaluating organizational compliance and security practices. Standards such as ISO 27001, NIST SP800-63, PCI DSS, and FIPS provide comprehensive details and requirements for information security, risk management, data protection, and privacy. Auditing against these standards helps organizations assess their adherence to best practices, identify gaps or vulnerabilities, and demonstrate their commitment to maintaining a secure and compliant environment.

Internal Standards

Organizations also establish internal standards to ensure the safety and integrity of operations and protect valuable resources such as data, intellectual property, and hardware. Internal standards provide consistent descriptions to define and manage important organizational practices. Standards differ from policies in a few ways. A simplistic view of the differences between the two is that standards focus on implementation, whereas policies focus on business practices.

Password standards describe the specific technical requirements required to design and implement systems, including how passwords are managed within those systems to ensure that different systems can interoperate and use consistent password-handling methods.

Hashing Algorithms —D efines requirements for the hash functions used to store passwords.

Password Salting —Defines the methods used to protect password hashes to protect them from rainbow table attacks.

Secure Password Transmission —Defines the methods for secure password transmission, including details regarding appropriate cipher suites.

Password Reset —Defines appropriate identity verification methods to protect password reset requests from exploitation.

Password Managers —Defines the requirements for password managers that organizations may choose to incorporate.

Access control standards ensure that only authorized individuals can access the systems and data they need to do their jobs to protect sensitive information and help prevent accidental changes or damage. Internally developed access control standards typically include the following elements:

Access Control Models —Defines appropriate access models for different use cases. Examples include role-based access c ontrol (RBAC), discretionary access control (DAC), and mandatory access control (MAC), among others.

User Identity Verification —Defines acceptable methods to verify identities before granting access. Examples include simple passwords, security tokens, biometric data, and other methods.

Privilege Management —Defines the methods for managing user privileges to ensure they have the minimum required access.

Authentication Protocols —Defines specific acceptable authentication protocols, such as Kerberos, OAuth, or SAML.

Session Management —Defines allowable session management practices, including requirements for session timeouts, secure generation and transmission of session cookies, and other similar requirements.

Audit Trails —Defines mandatory audit capabilities designed to assist with identifying and investigating security incidents.

Physical security standards protect datacenters, computer rooms, wiring closets, cabling, hardware, and infrastructure comprising the IT environment and the people who use and maintain them. Some examples include the following :

Building Security —Methods for securing facilities, including card access systems, CCTV surveillance, and security personnel.

Workstation Security —Standards for physically securing laptops or other portable devices.

Datacenter and Server Room Security —Define s requirements for card access, biometric scans, sign-in/sign-out logs, and escorted access for visitors.

Equipment Disposal —Defines requirements for securely disposing (or repurposing) equipment to ensure that sensitive data is irrecoverable.

Visitor Management —Defines the requirements for managing visitors, such as sign-in/sign-out procedures, visitor badges, and escorted access requirements.

Encryption protects data from unauthorized access, and it is vital for securing data both at rest (stored data) and in transit (data being transmitted). Encryption standards identify the acceptable cipher suites and expected procedures needed to provide assurance that data remains protected.

Encryption Algorithms —Defines allowable encryption algorithms, such as AES (Advanced Encryption Standard) for symmetric or ECC for asymmetric encryption.

Key Length —Defines the minimum allowable key lengths for different types of encryption.

Key Management —Defines how keys are generated, distributed, stored, and changed. It often includes requirements for using secure key management systems, procedures for regularly changing keys, and procedures for revoking them if they are compromised.

Procedures

Policies and guidelines set a framework for behavior. Procedures define step-by-step instructions and checklists for ensuring that a task is completed in a way that complies with policy.

Playbooks

Playbooks are essential to establishing and maintaining organizational procedures by establishing a central repository of well-defined, standardized strategies and tactics. They guide personnel to ensure consistency in operations and improve quality and effectiveness.

Playbooks facilitate knowledge sharing and continuity as employees move into new roles or leave the organization. Playbooks also mitigate risk by documenting critical procedures and preserving institutional knowledge. Playbooks help new team members quickly learn established processes while existing team members have a reference point for their tasks.

Moreover, playbooks act as a tool for quality assurance and continuous improvement. Clearly defining processes and the best practices to handle them makes it easier to identify and improve problem areas. By using playbooks, organizations can monitor the use and effectiveness of procedures over time and modify them as necessary to foster an environment of continual learning and development.

Most significantly, playbooks are essential in incident response and crisis management because they detail emergency procedures and contingency plans vital to steering activities during an emergency or crisis. Playbooks help incident response teams make quick decisions and work more effectively under stress, leading to more resilient operations and reducing the likelihood and impact of major security incidents.

Several best practice guides and frameworks are available to assist in developing playbooks, such as The MITRE ATT&CK framework https://attack.mitre.org , NIST Special Publication 800-61 https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final , and Open Source Security Automation (OSSA) https://www.opensecurityandsafetyalliance.org/About-Us .

11.1.4 Apply Appropriate Polices and Regulations

11.1.5 Practice Questions (Section Quiz)

q_proj_mgmt_oversight_ccpa_secp8

You are a data privacy officer for a tech company based in the United States. The company is planning to expand its services to California.

Which of the following legal regulations should be your primary focus for ensuring data privacy compliance?

Answers:

General Data Protection Regulation (GDPR)

Health Insurance Portability and Accountability Act (HIPAA)

***California Consumer Privacy Act (CCPA)**

Federal Information Security Management Act (FISMA)

Explanation:

California Consumer Privacy Act (CCPA) is the correct answer. The California Consumer Privacy Act (CCPA) gives California residents more control over the personal information that businesses collect about them. If the company is expanding its services to California, it must ensure compliance with the CCPA to avoid potential legal issues.

The General Data Protection Regulation (GDPR) applies to the processing of personal data of individuals who are in the European Union. While it's a comprehensive and strict regulation, it would not be the primary focus for a company expanding its services to California.

The Health Insurance Portability and Accountability Act (HIPAA) applies to entities that handle protected health information. Unless the tech company is dealing with health information, this would not be the primary focus.

The Federal Information Security Management Act (FISMA) applies to federal agencies and companies that contract with the government. Unless the tech company is dealing with government contracts, this would not be the primary focus.

q_proj_mgmt_oversight_centralized_secp8

An organization is restructuring its IT governance framework to improve its cybersecurity strategy. The organization has several distributed offices across various geographical regions, each having a unique set of IT policies and infrastructure.

The cybersecurity lead aims to increase control and consistency over the security practices in each office while retaining some autonomy for the individual offices to manage their specific risks.

Which governance structure aligns with the objectives of the cybersecurity lead and effectively mitigates risks associated with the security practices at each office?

Answers:

***Centralized governance with an advisory board**

Decentralized governance with an executive committee

Change Control Board (CCB)

Centralized governance with a steering committee

Explanation:

Centralized governance involves standardizing IT policies and practices across the organization, increasing control and consistency, which is the main aim of the cybersecurity lead in the given scenario.

Decentralized governance would allow each office to independently manage its IT policies and practices, which would not help the organization achieve control and consistency over security practices.

CCB primarily deals with assessing, approving, or rejecting proposed changes to a project or system. While a CCB can be a component of IT governance, it does not serve as a comprehensive governance structure.

Steering committees generally make strategic decisions at a higher level and may not be as involved in addressing specific issues the way an advisory board would be.

q_proj_mgmt_oversight_committee_secp8

Which of the following BEST describes a governance committee in the context of security governance?

Answers:

A group of high-ranking executives responsible for setting the strategic direction and policies of the organization.

***A specialized group comprised of subject matter experts, stakeholders, and representatives from relevant departments that focus on specific issues, such as security, risk management, audit, or compliance.**

A group of external stakeholders responsible for evaluating the security program's overall effectiveness.

A group of individuals responsible for the allocation of security resources and decision-making based on localized needs and priorities.

Explanation:

A governance committee is a specialized group that focuses on specific issues and provides in-depth analysis, recommendations, and operational support to the governance board.

A group of high-ranking executives responsible for setting the strategic direction and policies of the organization is incorrect because it describes the role of a governance board, not a governance committee. While both are important to an organization's governance structure, they serve distinct roles.

A group of external stakeholders responsible for evaluating the security program's overall effectiveness is incorrect because it describes a part of the role of a governance board, not a governance committee. Governance boards are responsible for evaluating the security program's overall effectiveness, among other duties.

A group of individuals responsible for the allocation of security resources and decision-making based on localized needs and priorities is incorrect because it describes a characteristic of decentralized security governance, not a governance committee. In a decentralized model, decision-making authority and resource allocation are distributed to different groups or departments.

q_proj_mgmt_oversight_decentralized_secp8

A nationwide company realizes its current standardized approach to security is not working. The different company business units need more autonomy and the ability to make decisions that meet their local needs and priorities.

What type of security governance should they follow?

Answers:

***Decentralized security governance**

Centralized security governance

Governance committees

Data protection authorities

Explanation:

Decentralized security governance distributes decision-making authority to different groups or departments to facilitate security-focused decisions based on localized needs and priorities.

In centralized security governance, decision-making authority primarily rests with a single core group or department that establishes policies, procedures, and guidelines and makes important security-focused decisions.

Specialized groups, such as subject matter experts, stakeholders, and representatives from relevant departments, comprise governance committees that focus on specific issues, such as security, risk management, audit, or compliance.

Data protection authorities focus on protecting personal data and privacy rights. They enforce data protection regulations and provide guidance on the best practices for securing personal information.

q_proj_mgmt_oversight_gdpr_secp8

Under the General Data Protection Regulation (GDPR), which of the following statements is correct regarding the rights of data subjects?

Answers:

***Data subjects have the right to be forgotten, but this is limited to information that is no longer necessary for the purpose it was collected.**

Data subjects have the right to data portability, but only within the same industry.

Data subjects have the right to access their personal data, but they must pay a fee for this service.

Data subjects have the right to rectification, but only within 30 days of data collection.

Explanation:

Data subjects have the right to be forgotten, but this is limited to information that is no longer necessary for the purpose it was collected is the correct answer. Under the GDPR, data subjects have the right to be forgotten, also known as the right to erasure. This means they can request the deletion or removal of personal data where there is no compelling reason for its continued processing. This includes situations where the data is no longer necessary in relation to the purpose for which it was originally collected or processed.

The right to data portability allows data subjects to obtain and reuse their personal data across different services, not just within the same industry. They have the right to move, copy or transfer personal data easily from one IT environment to another in a safe and secure way.

Under the GDPR, data subjects have the right to access their personal data free of charge. This includes the right to obtain a copy of their personal data, as well as other supplementary information.

The right to rectification means that data subjects can have their personal data corrected if it is inaccurate or incomplete. There is no specific time limit, such as 30 days, for this right under the GDPR.

q_proj_mgmt_oversight_global_law_secp8

As a cybersecurity professional, you are working for a multinational company that operates in multiple countries. The company is planning to launch a new product that collects user data for personalization features.

Which of the following legal considerations should be your primary focus?

Answers:

The local laws of the country where the company is headquartered.

The laws of the country where the majority of users are located.

The laws of the country where the data servers are located.

***The global laws that apply to data privacy and protection.**

Explanation:

The global laws that apply to data privacy and protection is the correct answer. As a multinational company, it's crucial to consider global laws that apply to data privacy and protection. This ensures that the company is compliant with all relevant laws, regardless of where its users are located or where its operations are based. It also helps to prevent potential legal issues that could arise from differences in data privacy laws between countries.

While the local laws of the country where the company is headquartered are important, they may not cover all the legal requirements for data privacy and protection in other countries where the company operates or where its users are located.

Focusing solely on the laws of the country where the majority of users are located may leave the company exposed to legal risks in other countries. It's important to consider the laws of all countries where the company has users.

While the laws of the country where the data servers are located are important, they may not fully cover the legal requirements for data privacy and protection in other countries where the company operates or where its users are located.

q_proj_mgmt_oversight_governance_committee_secp8

A multinational company discovered its existing cybersecurity policies were no longer adequate due to evolving cybersecurity threats and updated industry regulations.

The board of directors, comprising high-ranking executives, decided to review and revise the policies.

Who should the company involve in this process?

Answers:

Data custodian

Data processor

***Governance committee**

Regulatory agency

Explanation:

A governance committee is a specialized group comprised of subject matter experts, stakeholders, and representatives from relevant departments and focuses on specific issues such as security, risk management, audit, or compliance.

While a data custodian, or data steward, is responsible for the safe custody, transport, and storage of data, this role is more about the implementation of business rules rather than the revision of high-level policies.

The role of a data processor is mainly to process personal data on behalf of the controller.

While the company must consider the regulations set by regulatory agencies when revising its policies, it is not typically necessary to directly involve these agencies in the process.

q_proj_mgmt_oversight_guidelines_secp8

Which of the following is NOT a guideline for implementing effective governance controls?

Answers:

Implement a governance structure that best supports the organization's objectives.

Leverage expertise through committees to support decision-making.

Establish a comprehensive list of policies, processes, standards, and guidelines.

***Avoid using automation and orchestration tools to maintain control and promote transparency.**

Explanation:

Avoiding using automation and orchestration tools to maintain control and promote transparency is NOT a guideline. The guideline is to use automation and orchestration tools to improve consistency, reduce response times, and support compliance. Avoiding these tools would not be a guideline for implementing effective governance controls.

Implementing a governance structure that best supports the organization's objectives is indeed a guideline for implementing effective governance controls. It ensures that the governance structure aligns with the organization's goals and strategies.

Leveraging expertise through committees to support decision-making is a guideline for implementing effective governance controls. Committees provide a platform for subject matter experts to contribute their knowledge and insights, thereby enhancing the decision-making process.

Establishing a comprehensive list of policies, processes, standards, and guidelines is a guideline for implementing effective governance controls. These elements provide a clear framework for the organization's operations and ensure consistency and compliance.

q_proj_mgmt_oversight_hipaa_secp8

Which of the following cybersecurity laws and regulations is specifically designed to protect the privacy and security of health information in the healthcare industry?

Answers:

The Sarbanes-Oxley Act (SOX)

The Payment Card Industry Data Security Standard (PCI DSS)

***The Health Insurance Portability and Accountability Act (HIPAA)**

The Federal Information Security Management Act (FISMA)

Explanation:

The Health Insurance Portability and Accountability Act (HIPAA) is the correct answer. HIPAA is a federal law that required the creation of national standards to protect sensitive patient health information from being disclosed without the patient's consent or knowledge. It is specifically designed for the healthcare industry.

The Sarbanes-Oxley Act (SOX) is a law that sets requirements for all US public company boards, management, and public accounting firms. It does not specifically focus on the healthcare industry or the protection of health information.

The Payment Card Industry Data Security Standard (PCI DSS) is a set of security standards designed to ensure that all companies that accept, process, store or transmit credit card information maintain a secure environment. It does not specifically focus on the healthcare industry or the protection of health information.

The Federal Information Security Management Act (FISMA) is a United States federal law enacted in 2002 as Title III of the E-Government Act of 2002. The act recognized the importance of information security to the economic and national security interests of the United States. It does not specifically focus on the healthcare industry or the protection of health information.

q_proj_mgmt_oversight_processor_secp8

Your organization is implementing a new data governance model. You are tasked with assigning roles to various team members.

One of your colleagues is responsible for processing personal data on behalf of the controller, based on the controller's instructions.

Which data governance role BEST fits this colleague's responsibilities?

Answers:

Owner

Controller

***Processor**

Custodian

Explanation:

Processor is this is the correct answer. The processor processes personal data on behalf of the controller, based on the controller's instructions. This aligns directly with the role described in the question.

The owner is a high-ranking employee, like a director or a vice president, who is ultimately responsible for ensuring data is appropriately protected. They identify what level of classification and sensitivity the data has, decide who should have access to it, and what level of security should be applied. This does not align with the role described in the question.

The controller determines why and how personal data is collected, stored, and used. While the controller does have a role in processing data, the role described in the question is specifically following the controller's instructions, which aligns more with the processor role.

The custodian is responsible for the day-to-day management and protection of data. While they may handle data, they do not process it on behalf of the controller, which is the role described in the question.

q_proj_mgmt_oversight_stakeholders_secp8

An IT manager prepares a proposal to implement change management. Before being able to start the program, the manager needs support from key personnel within every department.

What key personnel does the manager need support from?

Answers:

Controller

Owner

***Stakeholders**

Processor

Explanation:

Stakeholders in change management are personnel with a vested interest in the change. Their participation fosters ownership and responsibility.

The controller role closely relates to the General Data Protection Regulation (GDPR) and identifies the purposes, conditions, and means of processing personal data. This position is not part of change management.

A high-ranking employee, like a director or a vice president, typically holds the owner role and is ultimately responsible for ensuring appropriate data protection. This position is not part of change management.

The processor is responsible for processing personal data on behalf of the controller and often represents Cloud Service Providers (CSP). However, vendors and business partners could also be representatives. This position is not part of change management.

q_pol_standards_procedures_audits_secp8

In a healthcare organization, the IT department conducts regular internal audits to ensure compliance with Health Insurance Portability and Accountability (HIPAA) regulations and to identify potential vulnerabilities.

During an audit, the IT team gathers evidence to evaluate the effectiveness of the implemented security controls and assess the overall security posture. The IT team uses questionnaires to obtain information from various departments and employees as part of the audit process.

What is the purpose of gathering evidence from internal audits and questionnaires?

Answers:

***Evaluate security effectiveness and identify areas for improvement.**

Obtain employee feedback on the company's security measures.

Implement new security policies based on survey results.

Ensure compliance with financial regulations.

Explanation:

The IT department evaluates the effectiveness of security controls and identifies areas for improvement by gathering evidence from internal audits and questionnaires. This evidence allows the team to make informed decisions about enhancing security posture.

Questionnaires may collect employee feedback, but the primary purpose of gathering evidence from internal audits and questionnaires is to assess security effectiveness, not solely to obtain feedback.

The IT department uses the evidence collected to assess existing security measures rather than implementing new policies solely based on survey results.

Evidence gathered from internal audits and questionnaires serves a broader purpose than just compliance with financial regulations. It aims to evaluate security effectiveness and identify areas for improvement across the financial institution.

q_pol_standards_procedures_aup_secp8

A newly hired chief information security officer (CISO) is implementing the National Institute of Standards and Technology (NIST) Cybersecurity Framework.

What first function would help the CISO better develop the company's security policies, such as acceptable use policy (AUP), and build out recommendations for security controls?

Answers:

Protect

***Identify**

Detect

Respond

Explanation:

The identify function in the National Institute of Standards and Technology's Cybersecurity Framework refers to developing security policies and capabilities. The CISO preparing policies and controls would fall under the identify function.

The second function of the NIST Cybersecurity Framework is protect, referring to the procurement, development, or deployment of IT assets and how to defend them against malicious actors.

The framework's third function is detect, which refers to ongoing, proactive monitoring to ensure that controls are effective and capable of protecting against new threats.

The fourth function of the framework is respond, determining how a company would identify, analyze, contain, and eradicate threats to systems and data security.

q_pol_standards_procedures_coop_01_secp8

A large technology company has recently experienced a significant system failure due to a cyberattack. The chief information security officer (CISO) is conducting a post-incident review to identify ways to improve the organization's resilience and recovery capabilities.

The CISO wants to focus on strategies that could have prevented the system downtime or minimized its duration and impact.

From a resilience and recovery standpoint in security architecture and continuity of operations planning (COOP), which of the following strategies would the CISO MOST likely recommend implementing to enhance the organization's ability to prevent or quickly recover from similar incidents in the future? (Select two.)

Answers:

Expand the IT team with more developers

Invest in a stronger firewall system

***Establish a redundant data center**

***Implement a detailed incident response plan**

Establish strong password policies and standards

Explanation:

An organization can notably strengthen its resilience and recovery capabilities by establishing a redundant data center. The redundant site swiftly takes over operations in case of a system failure or cyberattack, reducing downtime.

Formulating a detailed incident response plan is key to guaranteeing an effective, coordinated response to incidents. This plan specifies the procedures for identifying, reacting to, and recovering from security incidents.

Expanding the IT team could help if the new employees simply focused on cyber security issues. However, it is also the most expensive solution, and more team members may not be as effective as re-organizing the existing team or giving them the resources and tools to take care of the issue.

Investing in a stronger firewall system does not always guarantee that system failures due to a cyberattack will be significantly reduced. In this scenario, it would be wiser to first establish a redundant data center and implementing a more detailed incident response plan. Then, if that is ineffective, strengthening the firewall might be a solution.

Strengthening password policies and standards can help to prevent or reduce the occurrence of cyberattacks. However, in this scenario, the CISO made the correct decision to first address the issue by establishing a redundant data center and implementing a more detailed incident response plan.

q_pol_standards_procedures_coop_02_secp8

A cybersecurity team plans to improve the resilience of their organization's IT infrastructure.

The lead architect suggests implementing continuity of operations planning (COOP) to address potential disruptions and keep critical operations running during unexpected events.

What primary objective BEST describes the purpose of implementing COOP within an organization's security architecture?

Answers:

***It ensures continuous critical operations during unexpected disruptions.**

It prevents unauthorized access to sensitive data.

It monitors network traffic for malicious activities.

It provides immediate power backups during electrical outages.

Explanation:

COOP focuses on maintaining the continuity of an organization's critical operations during unexpected disruptions, ensuring minimal downtime and impact.

While data protection is a security concept, it is not the primary objective of COOP. Other mechanisms like encryption and access control lists focus on this goal.

Network monitoring is essential for security, but it is not the main objective of COOP. This function typically falls under intrusion detection systems (IDS) or intrusion prevention systems (IPS).

While power backups are crucial, they form a part of business continuity planning (BCP) and Disaster Recovery (DR) more than COOP. COOP encompasses a broader range of operations beyond just power supply.

q_pol_standards_procedures_coop_03_secp8

An earthquake occurs near the company HQ, causing severe damage in the area. The earthquake affected the building, which will not be usable for several weeks.

What plan will the company follow to maintain its business? (Select two.)

Answers:

***COOP**

***Disaster recovery**

Incident response

AUP

Playbook

Explanation:

The continuity of operations plan (COOP) focuses on the critical processes that must remain operational during and after a substantial disruption, like a natural disaster.

A disaster recovery plan details the steps required to recover from a catastrophic event such as a natural disaster. The goal is to restore operations as quickly and efficiently as possible.

The following will not help the company maintain its business in this scenario:

The incident response policy outlines the processes to follow after a security breach or cyberattack occurs.

An acceptable use policy (AUP) outlines how individuals may use network and computer systems by defining what constitutes acceptable behavior by users.

Playbooks are essential to establishing and maintaining organizational procedures by establishing a central repository of well-defined, standardized strategies and tactics. They guide personnel to ensure consistency in operations and improve quality and effectiveness.

q_pol_standards_procedures_coop_04_secp8

A nonprofit organization is working to create an integrated strategy that responds to potential disasters and ensures the continuation of essential functions across various scenarios, including budget constraints and prolonged disruptions.

Which approach would BEST address these multifaceted requirements?

Answers:

***Establish a COOP.**

Deploy a cold site.

Set up a warm site.

Implement a hot site.

Explanation:

Establishing a continuity of operations plan (COOP) offers a holistic approach that considers disaster recovery and the continuation of essential functions across different challenges.

Deploying a cold site is a specific strategy focusing on cost-effective recovery after a disaster but does not cover the comprehensive requirements for ensuring essential functions across various scenarios.

Setting up a warm site gears toward quicker recovery than a cold site but lacks the broad organizational resilience needed to address multiple scenarios, including prolonged disruptions.

Implementing a hot site centers on immediate recovery but falls short of providing a complete solution for the organization's wider continuity needs and budget considerations.

q_pol_standards_procedures_data_privacy_secp8

A company is reviewing its policies to ensure compliance with data privacy regulations. It wants to establish a policy that outlines the appropriate use of customer data.

What type of policy should the company focus on?

Answers:

***Data privacy policy**

Acceptable use policy (AUP)

Information security policy

Disaster recovery policy

Explanation:

The company is specifically concerned about compliance with data privacy regulations. Therefore, it should develop a data privacy policy outlining customer data handling to maintain compliance with relevant laws and data privacy protection.

An AUP focuses on defining acceptable behavior by users regarding network and computer systems; it may not specifically address data privacy.

An information security policy is broader and covers guidelines related to information security within the organization, including data privacy. However, this scenario needs a more specific data privacy policy.

Unrelated to data privacy, a disaster recovery policy focuses on recovering from catastrophic events or disruptions.

q_pol_standards_procedures_incident_response_secp8

A company merged with another company and is reviewing and combining both companies' procedures for incident response.

What should the joined companies have at the end of this preparation phase?

Answers:

***Incident response plan**

Communication plan

Playbook

Incident response lifecycle

Explanation:

An incident response plan lists the procedures, contacts, and resources for various incident categories available to responders. It covers expected or past incidents and the appropriate steps to respond accordingly.

A communication plan is part of the incident response plan that establishes clear lines of communication for reporting incidents and notifying affected parties as the management of an incident progresses.

A playbook is a data-driven standard operating procedure assisting analysts in detecting and responding to specific cyber threat scenarios. It is part of an incident response plan.

The incident response lifecycle is a seven-step process under which the preparation stage falls. It is not a document presented during the preparation phase.

q_pol_standards_procedures_iso_iec_secp8

A large multinational company adopts a new standard to enhance its information security management system. The company operates across different regions, so the chosen standard must be internationally recognized.

The company wants the standard to provide a comprehensive framework to ensure adequate and proportionate security controls.

Which of the following standards would be MOST suitable for the company's needs?

Answers:

***ISO/IEC 27001**

ISO/IEC 27018

NIST Special Publication 800-63

PCI DSS

Explanation:

ISO/IEC 27001 provides a comprehensive framework for an information security management system (ISMS), ensuring adequate and proportionate security controls. It is suitable for international use and ideal for a multinational company.

The ISO/IEC 27018 standard also pertains to information security, protecting personally identifiable information (PII) in public clouds. Although it may also be useful, it is less comprehensive than ISO/IEC 27001 for general information security management.

The NIST Special Publication 800-63 standard is a U.S. government standard for digital identity guidelines. While it may offer useful guidelines for parts of the company's security needs, it is less comprehensive than ISO/IEC 27001.

Payment Card Industry Data Security Standard (PCI DSS) is more specific for organizations that handle credit card transactions from major card providers.

q_pol_standards_procedures_password_salting_secp8

As a security analyst at a large corporation, you are tasked with reviewing and improving the company's password security measures.

Currently, the company uses a simple hashing algorithm to store passwords. You are considering four options to enhance password security.

Which of the following would be the MOST effective method to implement?

Answers:

Increase the minimum password length requirement

Implement password complexity requirements

***Implement password salting**

Implement a password expiration policy

Explanation:

Implementing password salting is the most effective method among the options. Salting involves adding a unique random string of characters to each password before it is hashed. This makes pre-computed hash attacks, such as rainbow table attacks, ineffective and significantly increases the difficulty of cracking the passwords, even if the hash values are somehow compromised.

Increasing the minimum password length requirement can improve security by making brute force attacks more difficult. However, it does not protect against other types of attacks, such as rainbow table attacks, and can lead to user frustration and poor password practices, such as writing down passwords.

Implementing password complexity requirements can also deter brute force attacks and make guessing passwords more difficult. However, like increasing password length, it does not protect against all types of attacks and can lead to poor password practices.

Implementing a password expiration policy can help limit the potential damage if a password is compromised. However, it can also lead to user frustration and poor password practices, such as making minor modifications to old passwords or writing down passwords. Furthermore, if a password is compromised, the attacker usually exploits it immediately, making the expiration policy less effective.

q_pol_standards_procedures_playbooks_secp8

A company's security team prepares a playbook as part of its business continuity planning. The playbook outlines various scenarios, including natural disasters, cyberattacks, and power outages, detailing how to minimize business interruption and swiftly restore systems.

Given the scenario, what is the primary function of the playbook within the company's business continuity plan?

Answers:

***Directing immediate response and recovery operations**

Predicting the occurrence of future disaster events

Eliminating the risk of natural disasters and cyberattacks

Training employees on how to prevent power outages

Explanation:

The playbook's role is to guide immediate response and recovery actions during a disruption. Playbooks provide step-by-step instructions for the security team, ensuring a swift, coordinated response to various scenarios outlined in the business continuity plan.

While playbooks include potential scenarios, they do not predict when these events will occur. Their purpose is to guide responses when such events happen.

Playbooks do not eliminate risks but prepare teams to manage and respond effectively.

While employee training can be a component of a comprehensive business continuity plan, it is not the primary purpose of a playbook. Playbooks manage responses during events, not prevent them.

q_pol_standards_procedures_session_management_secp8

As a security manager at a financial institution, you are reviewing the company's access control measures. You have identified potential areas of improvement and are considering four options to enhance access control security.

Which of the following would be the MOST effective method to implement?

Answers:

Implement biometric authentication

Increase password complexity requirements

***Improve session management**

Implement a two-factor authentication

Explanation:

Improving session management is the most effective method among the options. Good session management includes measures like session timeout, which logs a user out after a period of inactivity, and session invalidation, which ends the session when the task is completed or when the user logs out. This can prevent unauthorized access if a user leaves their device unattended or if a session token is compromised.

Implementing biometric authentication can enhance security by adding a layer of physical identity verification. However, it can be costly to implement, and there are potential privacy concerns. Moreover, it does not address issues such as session hijacking, where an attacker can gain control after authentication has occurred.

Increasing password complexity requirements can deter brute force attacks and make guessing passwords more difficult. However, it does not protect against attacks that occur after authentication, such as session hijacking or cross-site scripting attacks.

Implementing two-factor authentication can significantly improve security by requiring a second form of verification in addition to a password. However, like biometric authentication and password complexity requirements, it does not address attacks that occur after authentication.

11.2 Change Management

As you study this section, answer the following questions:

When is change management used, and what risks does it help mitigate?

Who should be involved in change management?

How do allow and deny lists help in change management?

What is the purpose of version control in change management?

The key terms for this section include:

Term	Definition
Stakeholders	A person who has a business interest in the outcome of a project or is actively involved in its work.
Dependencies	Resources and other services that must be available and running for a service to start.
Version control	The practice of ensuring that the assets that make up a project are closely managed when it comes time to make changes.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	1.3 Explain the importance of change management processes and the impact to security. Technical implications Allow lists/deny lists Restricted activities Downtime Service restart

	Application restart Legacy applications Dependencies Documentation Updating diagrams Updating policies/procedures Version control 5.1 Summarize elements of effective security governance. Policies Change management Procedures Change management Types of governance structures Boards 5.2 Explain elements of the risk management process. Risk analysis

11.2.1 Change Management (Lesson Video)

Transcript:

Change is inevitable in the IT field. Whether it's updating to new equipment or changing direction mid project, you're bound to come across a lot of changes!

In this video, we're going to talk about change management. This is just the method we use to ensure that changes are implemented as smoothly and effectively as possible.
Most organizations have an official approval process for change management. A business commonly uses change order forms to document these requests.
Change order forms are different for each company, but they typically ask for the change's purpose, scope, and extent, along with a risk analysis and rollback plan in case things don't go as expected.
So, after filling out your change order form, what's next? Well, first, it must be approved by the change board. A change board might consist of one person or many people who approve changes to the project. Board members might be managers, project supervisors, and maybe even the CEO. If you were developing this project for a customer instead of for in-house use, the customer would need to approve the changes as well.

The first thing the change board will want to know is why you want to make this change at all. You'll want to explain why the change is needed and whether it'll improve a current process, save money, or add additional functionality. You might also be asked to explain how it fits within the organization's current strategic focus. The change board will also want to know the scopeâ€”what exactly the proposed change will entail from start to finish. A scope has many components, but it should always include outcomes, cost, personnel, and systems.

Some other questions you'll need to have a ready answer for are: What are your desired outcomes? What exactly do you want to change? What are the desired deliverables? How will you know when the project is complete? The board will also obviously want to know how much this is all going to cost, whether in money, time, or resources.

You'll want to consider which personnel you'll need to implement the change. Will you need developers, technicians, or administrative staff? Be sure to include the contact information for a point person who can serve as the change board's primary contact for the project, if it's approved. It's also important to explain which existing systems will be impacted. Most likely, you'll be working alongside live production systems and any changes should be carefully planned out. Will there be downtime or potential conflicts? If so, you probably want to set up a sandbox to test the proposed changes before you implement them. As you can see, if there's literally anything you think you might need for your project, you want to include it in the scope.

Before you propose a change, you also want to analyze the potential risk. Yes, even small projects come with some risk. Be sure to ask yourself, "What could go wrong if we implement these changes?" After you think through that, the next obvious questions are: How are we going to avoid these problems? How difficult will it be to backtrack if things turn out worse than expected? Is the outcome worth the risk? If it turns out that the risks outweigh the rewards, you probably want to rethink your approach.

Now, before you start anything, make sure you get written approval for your changes. Written acceptance protects you, the change board, and the customer. This ensures that all of you have a complete understanding of the proposed changes and what they could entail. You want to avoid misunderstandings that could hurt your own or your organization's reputation or that might end up causing expensive problems.

Documentation is critical to any project from start to finish! During the implementation stage, you want to be sure to document all the changes you've made along the way. Include the specific approvals you've received. Once you're done, be sure to follow your organization's end user acceptance process. This ensures that someone who's going to be using the system verifies that it's working as expected. And be sure to document this as well.

That's it for this lesson. In this lesson, we discussed change management. We reviewed the information that's commonly needed for a change request. Then we went over the importance of documentation during a proposed change's approval and implementation phase.

11.2.2 Change Management Facts

Change management is a systematic approach to managing all changes made within an IT infrastructure. Its primary goal is to minimize risk and disruption while maximizing the value and efficiency of organizational changes. Change management relies on effective planning, testing, approval, and implementation of various changes ranging from minor updates to complicated system migrations. Change management must consider the potential impacts and dependencies of all proposed changes and mandate the development of contingency and rollback plans in case changes lead to unforeseen problems. Proper documentation and communication are also essential, ensuring that all relevant stakeholders understand the details of proposed changes and their implications. Change management programs promote a controlled progression of software and the IT infrastructure and contribute to its resilience and stability.

This lesson covers the following topics:

- Change management programs
- Factors driving change management
- Allowed and blocked changes
- Restarts, dependencies, and downtime
- Legacy systems and applications

Documentation and version control

Change Management Programs

Change management plays a vital role in an organization's security operations. It refers to a systematic approach that manages all changes made to a product or system, ensuring that methods and procedures are used to handle these changes efficiently and effectively. This helps minimize risks associated with the changes, ensuring they do not negatively impact the organization's security posture, service availability, or performance.

A non-comprehensive list of changes typically managed in a change management program includes the following:

- Software deployments
- System updates
- Software patching
- Hardware replacements or upgrades
- Network modifications
- Changes to system configurations
- New product implementations
- New software integrations
- Changes and refreshes to support environments

If not properly managed, these changes can introduce new vulnerabilities into the system, disrupt services, or negatively impact the organization's compliance status. A robust change management program allows all changes to be tracked, assessed, approved, and reviewed. Each change must include documentation, including details describing what will be changed, the reasons for the change, any potential impacts, and a rollback plan in case the change does not work as planned. Each change must be subject to risk assessment to identify potential security impacts. Appropriate personnel must approve changes before implementation to ensure accountability and ensure changes align with business priorities.

After implementations, changes must be reviewed and audited to ensure they have been completed correctly and achieved their stated outcome without compromising security. Systematic management of changes supports an organization's ability to reduce unexpected downtime and system vulnerabilities. Change management programs contribute to operational resilience by ensuring that changes support business objectives without compromising security or compliance.

A typical change management approval process involves several stages designed to ensure proper assessment and approval of change proposals. Change requests usually begin with submitting a request for change (RFC) that outlines the details of the proposed change, including its purpose, scope, and potential impact. The change request is reviewed by a designated change manager or committee that assesses its feasibility, risks, alignment with organizational objectives, and policy compliance. Following initial review, the change request undergoes a formal approval process involving relevant stakeholders, such as management, IT teams, and any impacted departments, to ensure consensus and authorization before the change is implemented. Throughout the process, documentation and communication are crucial in tracking the status and outcome of approved changes.

The implementation of changes should be carefully planned, with consideration for how the change will affect dependent components. For most significant or major changes, organizations should attempt to trial the change first. Every change should be accompanied by a rollback (or remediation) plan so that the change can be reversed if it has harmful or unforeseen

consequences. Changes should also be scheduled sensitively if they are likely to cause system downtime or other negative impacts on the workflow of the business units that depend on the IT system being modified. Most networks have a scheduled maintenance window period for authorized downtime. When the change has been implemented, its impact should be assessed, and the process reviewed and documented to identify any outcomes that could help future change management projects.

Factors Driving Change Management

Change management requires the expertise of individuals from various parts of an organization to oversee and implement changes effectively. Examples include IT professionals with technical knowledge, business leaders with operational knowledge, and compliance officers with legal expertise. The involvement of these stakeholders (which includes anyone with a vested interest in the change or project being implemented or developed) facilitates a comprehensive review of proposed changes, helping to identify non-obvious risks and identify effective implementation plans that minimize risks and business disruptions. Additionally, including diverse stakeholders promotes acceptance and adoption of the changes because they were involved in the planning and decision-making process. Stakeholder participation fosters ownership and responsibility, which are crucial for successful change implementation.

Ownership in change management refers to individuals or groups primarily responsible for implementing a specific change. Owners can be project managers, team leaders, or anyone responsible for the change. Owners are accountable for ensuring that the change is implemented as planned, risks are managed effectively, and there's a clear plan for communication and training associated with the change. They also ensure that all stakeholders appropriately review and approve the proposed change. Stakeholders in the change management process describe the individuals or groups impacted (or interested) in the change and include employees, managers, the Change Advisory Board (CAB), and sometimes even customers, vendors, and partners. Stakeholder engagement is critical to successful change management because keeping stakeholders informed about changes, understanding their concerns, and addressing their needs improves the likelihood of the change being accepted and implemented smoothly.

Allowed and Blocked Changes

Allow lists and deny lists play a significant role in change management practices and can reference two different viewpoints. One scenario views allow and deny lists in relation to change types or from a different view. Allow and block lists describe software restriction approaches designed to control computer software. In terms of change management, an allow list describes a list of approved software, hardware, and specific change types (such as routine or low-risk changes) that are not required to go through the entire change management process. An allow list may also include specific individuals with change management approval authority. Allow lists help streamline change management by reducing the time and effort required for trusted or preauthorized changes. Allow lists must be updated via regular reviews to stay corrected with changing organizational needs.

A deny list includes explicitly blocked software, hardware, and specific change types. The block list might include software and hardware with known security or compatibility issues, high-risk or high-impact changes that must always go through the full change management process, or individuals who are not authorized to implement or approve changes. In this regard, deny lists help prevent unauthorized or risky changes from being implemented. They can serve as a security measure to clearly identify off-limits change types so there is no room for negotiation or misinterpretation.

Allow and block lists also refer to technical controls that exist in a few different contexts, including access controls, firewall rules, and software restriction mechanisms. Allow and block lists can impact change implementation by causing unintended problems. For example, software allow lists can be negatively impacted by software patching. If allow lists are based on executable file hash values, they will fail to recognize newly patched executables after patching because their hash values will change. This can result in fully patched systems that are unusable by employees because none of the previously allowed software can run. Regarding change management, it is important to incorporate the potential impacts of allow and block lists into the testing plan.

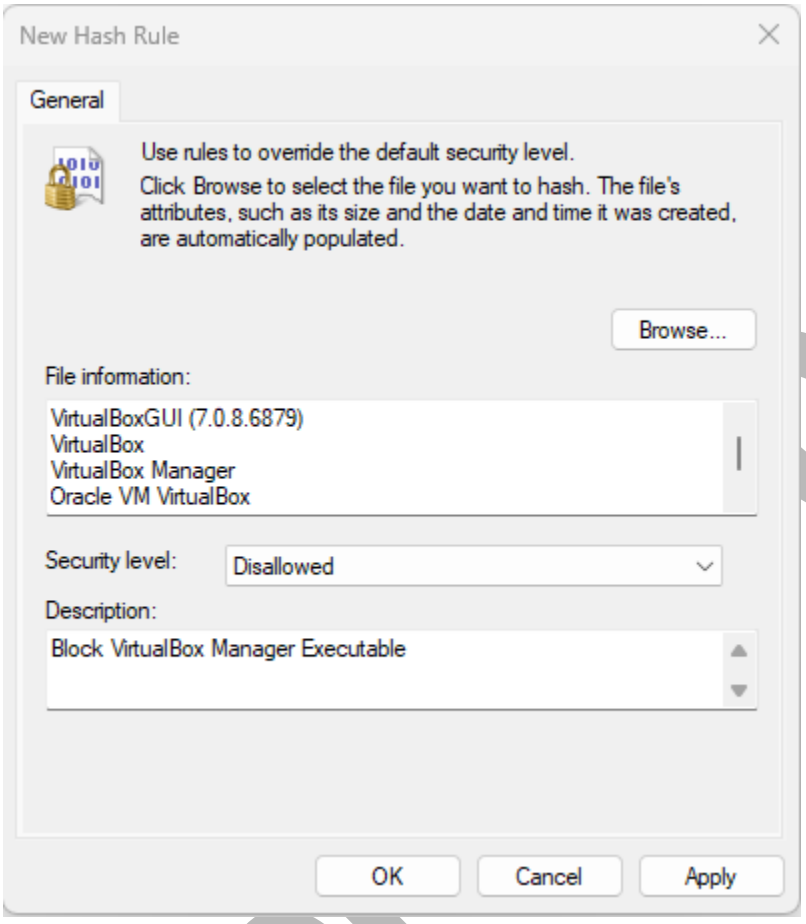

Software Restriction Policies (block list) can be based on file hash values. (Screenshot used with permission from Microsoft.)

Restricted activities refer to actions or changes that require additional scrutiny, strict controls, or higher levels of approval/authorization due to their potential impact on critical systems, sensitive data, or regulatory compliance.

Restarts, Dependencies, and Downtime

Service and application restarts, as well as downtime, are critical considerations because they typically have a direct impact on business operations. For example, reconfigurations and patching changes often require restarting services or applications, leading to downtime. One of the primary goals of change management is to minimize these disruptions by scheduling restarts or downtime events during maintenance windows or off-peak times to reduce the impact on users and business processes.

Change management processes include communication requirements designed to ensure relevant stakeholders are aware of service outages so they can prepare accordingly. Effective change communication enhances the visibility of the change management process among stakeholders and fosters a culture of transparency and cooperation.

Services and applications often depend on other software, interfaces, and services to function correctly. These dependencies complicate changes because a service restart in one area may significantly impact another. For example, if a database server is restarted, all applications that rely (depend) on the database will likely experience issues or downtime. A change that initially appeared to be minor may impact a wide range of the organization's operations. A careful analysis of software and system dependencies is critical for reasons like these. Understanding what services depend on each other, how restarts impact them, and what measures need to be taken to mitigate potential impacts help avoid unintended outages.

Dependencies also impact the time needed for a change. If a service restart requires other services to be shut down or restarted, the overall change process will need more time. Additionally, backout plans may also need to consider dependencies as part of the process, which will also require additional time.

Understanding the risks associated with restarts and downtime drives the development of effective backout plans and downtime contingencies, ensuring the organization is well-prepared to handle any potential complications and unintended consequences related to the change. Additionally, understanding risks associated with a change also supports the development of post-change performance monitoring to validate that systems function as required and help detect issues quickly. Sometimes, the potential risks of a change causing significant disruption require the organization to identify alternative solutions.

Some typical IT changes that generally require service or application restarts and result in downtime are as follows:

Change Type	Description
Software upgrades and patches	When upgrading software applications, especially major version updates or patches, a restart of the application is typically needed to apply the changes effectively and ensure the updated version is fully functional.
Configuration changes	Many system configuration changes, such as modifying server settings, network configurations, or database parameters, require a restart of the affected services to apply the changed configurations properly.
Infrastructure changes	When changing infrastructure components, such as switches, routers, firewalls, and load balancers, it is typically necessary to restart the devices to apply the changes and ensure they do not negatively impact operation.
Security changes	Implementing specific security measures, such as updating encryption protocols, enabling or disabling security features, or modifying access control settings, may require a restart of the services or applications to enforce the new security configurations effectively.

Downtime refers to the scheduled time designated for changes to be implemented (scheduled downtime) or the amount of time a service or application is unavailable (unscheduled downtime).

Legacy Systems and Applications

Legacy applications pose unique challenges regarding change management as these systems are often critical to business operations and are difficult to manage. Many legacy applications are built using outdated technology, which introduces compatibility issues when implementing changes. For example, new software or security updates may be incompatible with legacy systems. These Incompatibility issues might require specialized solutions, such as virtualization, emulation, interpreters, custom "fit-gap" software, or modifications to the newer components to ensure compatibility. These accommodations further complicate the manageability of the legacy application.

Legacy applications often lack comprehensive documentation or have been heavily customized over years or even decades, making them extremely difficult to manage. This complexity necessitates extensive testing and meticulous implementation plans to help avoid unintended consequences or outages. Legacy applications also typically lack vendor support, removing the option to "call for help," which increases the risks associated with any change. A lack of vendor support coupled with high complexity, poor documentation, and business criticality make legacy systems a significant security problem.

Documentation and Version Control

Version control refers to tracking and controlling changes to documents, code, or other important data. Organizations can use version control to maintain a historical record of changes, ensure only approved changes are implemented, and quickly revert changes to a previous version as warranted. Version control is also important when diagrams, policies, and procedures require updates. In this way, version control prevents confusion associated with using outdated or inconsistent documents.

Assessing how a change impacts existing policies, procedures, and diagrams is essential, and change management plans should include provisions requiring updates to these documents as part of the implementation. The frequency of diagram and documentation updates varies, but they are typically updated whenever significant changes or modifications to a process, system, or application occur. Once document updates have been completed, the new versions should be clearly labeled, and the older versions should be archived but still available for reference. Major changes may necessitate training for relevant teams or departments.

Change management is a crucial aspect of implementing changes to a system or application. By assessing the potential technical implications of these changes, organizations can take necessary steps to minimize disruptions. Effective change management requires following specific processes, such as developing implementation plans and conducting thorough testing procedures. Through change management, leadership can ensure that any changes made are successful and contribute positively to the organization.

Some examples of different documentation impacted by change management include the following:

Item	Description
Change requests	Change requests themselves should be reviewed and updated to reflect the details and status of the change, including any modifications or approvals during the change management process.
Policies and procedures	Changes may impact existing policies and procedures. As a result, these documents need to be reviewed and updated to ensure they align with the new processes, guidelines, or controls introduced through the change.
System or process documentation	Documentation should reflect any changes to systems, applications, or processes. It may involve updating system architecture, diagrams, process flows, standard operating procedures (SOPs), or user manuals to represent the current state and functionality of the changed system.
Configuration management documentation	Changes to configuration items, such as servers, networks, or databases, should be tracked and documented within the configuration management system to maintain an accurate record of its configuration.
Training materials	Changes often impact employees, and they may require more training. Existing training materials, such as presentations, manuals, or computer-based learning modules, must be reviewed and updated as warranted.
Incident response and recovery plans	Changes made to systems or applications may necessitate updates to incident response and recovery plans to ensure they account for the revised configurations, new dependencies, or recovery procedures resulting from the change.

Policies and procedures must change as often as technology does, which is often!

11.2.3 Practice Questions (Section Quiz)

q_change_management_allow_list_secp8

A recently breached company tasks the cyber team to further restrict end-user permissions.

What describes the use of an application allow list?

Answers:

***It enforces policies in computer systems and networks.**

It is used in computer systems and networks to enforce policies.

It is a list of rules or entries that specify users' access.

It controls access to files, directories, or systems resources in OSs.

Explanation:

An allow list (or approved list) will deny an execution unless it is a process that the organization explicitly authorizes.

Organizations can use access control lists (ACLs) in computer systems and networks to enforce access control policies.

An ACL is a list of rules or entries that specify a user or a group's access to specific resources or ability to perform certain actions.

ACLs control access to operating and file systems' files, directories, or system resources. Each access control entry (ACE) typically contains a user or group identifier and the associated permissions that allow or deny actions.

q_change_management_app_changes_secp8

At a medium-scale software development firm, significant modifications to several critical applications employees use daily are on the horizon.

Considering the principles of change management, what should the primary focus be during the implementation phase of these changes?

Answers:

***Scheduling service restarts during non-business hours to minimize application downtime.**

Ensuring high-risk alterations are prioritized and subject to the entire change management process.

Updating the block list to encompass previous versions of the applications.

Concentrating solely on the allow list while disregarding the block list for a smoother transition.

Explanation:

One of the main objectives of change management is to minimize disruptions to business operations. The company can achieve this by scheduling service restarts during off-peak hours or maintenance windows.

While high-risk alterations should indeed be subject to the entire change management process, this should not be the primary focus during the implementation phase of the changes.

While including previous versions of applications in the block list could be a part of change management, it should not be the primary focus while implementing the changes.

Solely concentrating on the allow list while disregarding the block list would not be advisable. Both lists play crucial roles in the change management process and implementation.

q_change_management_backout_plan_secp8

A software patch was inadvertently pushed out early, during the middle of the workday, and has brought business to a halt. The chief executive officer (CEO) demands that the systems return to full operations immediately.

What part of the change plan will assist in this task?

Answers:

***Backout plan**

Impact analysis

Test results

Standard operating procedures

Explanation:

A backout plan is for reversing changes and returning systems and software to their original state if the implementation plan fails. A well-defined backout plan helps to minimize downtime and reduces the risk of data loss or other severe impacts.

An impact analysis identifies and assesses the potential implications of a proposed change, including how the change will impact all areas.

Evaluation of changes in a test environment ensures they work as intended. Test results provide valuable insight into the likelihood of success without impacting business operations.

Standard operating procedures (SOPs) describe how to carry out routine operations or changes. In change management, SOPs ensure that change implementation is consistent and effective.

q_change_management_change_times_secp8

A properly implemented change plan for an international company helps keep business operations moving forward. Restarts, dependencies, and downtime are hand-in-hand with change management.

When is the BEST time to implement changes? (Select two.)

Answers:

After the work day

***Off-peak times**

Peak times

***Maintenance windows**

During holidays

Explanation:

If an organization needs to perform a change during the day, it should take place during off-peak times, broadcasting the intention to all the users.

Maintenance windows consist of scheduled times when services can safely restart to allow for all changes made per the change management policy.

While after the work day is possibly "off-peak time," many users work past a normal quitting time to focus on critical tasking.

Any changes should avoid peak times when users are at their most active work schedule, and implementing a change would cause the most disruption.

While holidays could be a good time to implement changes, in an international company holidays vary from location to location and its not always easy to find a consistent time to make the changes. In addition, the IT administrators making the changes would need to use holiday time to make the changes.

q_change_management_cmb_secp8

An educational institution plans to transition from a traditional to a digital learning system. The school's administration has assembled a Change Management Board (CMB) to ensure smooth and secure execution.

Within the scope of this educational institution's digital transformation project, what would be the primary role of the CMB?

Answers:

***Assessing, approving, and managing changes in the IT infrastructure**

Conducting cybersecurity awareness training for all school staff

Developing and implementing the school's digital learning curriculum

Overseeing the daily operations of the school's IT department

Explanation:

A CMB in any organization actively evaluates proposed changes, determines their potential impact, decides on their implementation, and manages these changes to minimize disruption and risk.

A security training team or department generally takes responsibility for this important aspect of overall cybersecurity, not the CMB.

Curriculum development and implementation are tasks that academic and educational technology departments actively perform. The CMB becomes involved only when these changes influence the IT infrastructure.

Managing changes to the IT infrastructure is the primary focus of the CMB, not overseeing the day-to-day operations of the IT department. IT managers or department heads usually take on this responsibility.

q_change_management_dependencies_secp8

An organization frequently implements changes, reconfigurations, and patches to enhance its IT infrastructure's security and efficiency. The cybersecurity analyst must carefully analyze dependencies between services, applications, and interfaces to avoid unintended outages and disruptions during service restarts or downtime events.

How does understanding dependencies impact the change management process? (Select the three best options.)

Answers:

***Helps avoid unintended outages and disruptions during service restarts or downtime events.**

***Guides the development of effective backout plans and downtime contingencies.**

***Supports the development of post-change performance monitoring to validate system functionality and quickly detect issues.**

Helps minimize the need for backout plans during changes.

Increases the involvement of stakeholders in the change management process.

Informs individuals or groups that are primarily responsible for implementing a specific change.

Guides the approval process to ensure proper assessment and approval of change proposals.

Explanation:

The following are areas for understanding how dependencies impact the change management process:

By knowing the dependencies, we can avoid unintended outages and disruptions during service restarts or downtime events and ensure that changes do not negatively impact interconnected services or applications.

When we understand the dependencies, we can guide the development of effective backout plans and downtime contingencies, thus preparing the organization to handle potential complications during changes.

By understanding the dependencies, we support the development of post-change performance monitoring to validate system functionality and promptly detect any issues that may arise after the change.

While understanding dependencies is essential for effective change management, it does not necessarily increase understanding in the following areas:

Minimizing the need for backout plans.

The involvement of stakeholders in the change management process.

Individuals or groups that are primarily responsible for implementing a specific change.

The approval process to ensure proper assessment and approval of change proposals.

q_change_management_legacy_apps_secp8

You are the IT manager at a large corporation. The company has been using a legacy application for several years. The application is critical for daily operations but it's not compatible with newer technologies the company plans to adopt.

The vendor no longer supports the application and it has known security vulnerabilities.

What should you do?

Answers:

Continue using the legacy application and hope no security breaches occur.

Immediately stop using the legacy application and switch to the new technology.

***Develop a plan to phase out the legacy application while adopting the new technology.**

Ignore the new technology and focus on finding a new vendor to support the legacy application.

Explanation:

Developing a plan to phase out the legacy application allows the company to address the security vulnerabilities and compatibility issues while minimizing disruption to daily operations is the best option. The plan should include risk assessment, training for users, data migration, and other key considerations.

Continuing to use the legacy application and hope no security breaches occur is not a good option because it leaves the company vulnerable to potential security breaches. Ignoring known vulnerabilities is not a responsible or effective way to manage IT risks.

Immediately stopping use of the legacy application and switch to the new technology could disrupt daily operations and cause significant problems. Switching to a new technology without a transition plan could lead to data loss, downtime, and other issues.

Ignoring the new technology and focusing on finding a new vendor to support the legacy application might seem like a solution but it's not practical or efficient. The legacy application is not compatible with newer technologies, and finding a new vendor to support outdated software could be difficult and costly. It also doesn't address the security vulnerabilities.

q_change_management_ownership_roles_secp8

After receiving the annual audit results from the Inspector General's office, a cyber specialist begins identifying points of contact to implement change management on numerous flagged processes.

Understanding the various positions tied to change management, which roles would normally have ownership in the change management process? (Select two.)

Answers:

***Project manager**

***Team leader**

Vendors

Partners

Human resources

Explanation:

Ownership in change management refers to individuals or groups primarily responsible for implementing a specific change and may also include project managers.

The team leader's role is also considered a big part of the ownership group. Implementing change as planned, managing risk effectively, and a clear plan for communication and training associated with the change are all part of accountability for this leadership role.

The role of a vendor is a stakeholder in the change management process, so there is no direct tie to ownership.

Similarly, the role of a partner is also a stakeholder, so there is no association with ownership responsibilities.

Human resources is not a stakeholder in the change management process, and would have no association with ownership responsibilities.

q_change_management_plan_change_components_secp8

Due to the introduction of security vulnerabilities during a previous change, company leadership wants reassurance that the vulnerabilities will not happen again.

The IT department has made several changes to its change management plan.

What are items the IT department would add to this plan? (Select three.)

Answers:

***Impact analysis**

***Test results**

***Backout plans**

Stakeholders

Rack layout diagram

Cable plan

Balance sheet

Explanation:

The following are items that the IT department would normally add to the plan:

Impact analysis identifies and assesses the potential implications of a proposed change, including how the change will impact all areas.

Evaluation of changes in a test environment ensures they work as intended. Test results provide valuable insight into the likelihood of success without impacting business operations.

A backout plan is for reversing changes and returning systems and software to their original state if the implementation plan fails. A well-defined backout plan helps to minimize downtime and reduces the risk of data loss or other severe impacts.

The following are items that the IT department would not normally add to the plan for this scenario:

Stakeholders are personnel with a vested interest in completing the change, but they are not components of a change management plan that would keep vulnerabilities from appearing.

A rack layout diagram and a cable plan are both networking documents for installing and configuring network hardware and are not directly related to correcting security vulnerabilities.

A balance sheet is a financial statement that reports a company's assets, liabilities, and shareholder equity. It is not directly related to a change management plan.

q_change_management_security_updates_secp8

Several mission-essential applications stopped working the morning after implementing a mandatory security update. As a result, implementing the security updates introduced several instabilities in existing software.

What could have prevented this from occurring?

Answers:

Stakeholders

***Change management**

Dependencies

Request for change

Explanation:

If the organization does not have a proper change management program for security updates, it can introduce new vulnerabilities into the system, disrupt services, or negatively impact its compliance status. It will also create additional work to roll everything back to its previous state.

Stakeholders are part of the change management program, including IT professionals with technical knowledge, business leaders with operational knowledge, and compliance officers with legal expertise.

Services and applications often have dependencies on other software, interfaces, and services to function correctly. If not planned correctly, changing one area can make all other applications have problems. This is another reason for a change management program.

A request for change is part of the change management program, and it would be an aspect of the solution, not the solution itself.

q_change_management_software_update_01_secp8

A financial institution receives a significant software update.

What is the optimal approach to handle this situation in a change management program?

Answers:

***Assess impact, test, get approval, apply update.**

Apply to critical systems first, then the rest.

Apply at next maintenance window without assessment.

Update systems with past vulnerabilities only.

Explanation:

In an effective change management program, it is crucial to assess the impact of the update, test it in a controlled environment, get the necessary approvals, and then apply the update system-wide. This step-by-step approach helps mitigate risks and ensure the update aligns with the organization's business objectives.

While applying the updates to a few critical systems first and then the rest seems cautious, it skips important steps such as assessment, testing, and approval.

Applying the update during the next maintenance window without prior assessment, testing, and approval could introduce unexpected risks.

Applying updates only to systems that have shown past vulnerabilities might leave other systems at risk and ignores the assessment process.

q_change_management_software_update_02_secp8

The organization is implementing a significant software upgrade that necessitates application restarts.

How can the cybersecurity analyst ensure a smooth transition without causing extended downtime?

Answers:

***Schedule the upgrade during nonworking hours to reduce the impact on users.**

Conduct the software upgrade without restarting the applications to avoid interruptions.

Restart all applications simultaneously to complete the upgrade faster.

Implement the upgrade without analyzing software dependencies.

Explanation:

The cybersecurity analyst can help minimize disruptions to users and business processes by scheduling the software upgrade during nonworking hours. This approach aligns with change management's goal to reduce the restart impact on business operations.

Restarting the applications is typically necessary to apply the software upgrade changes effectively.

If all the applications restart simultaneously, it can lead to extended downtime and further complicate the troubleshooting process in case of issues.

Implementing the upgrade without analyzing software dependencies is incorrect since understanding potential impacts and ensuring a successful upgrade is crucial.

q_change_management_sop_01_secp8

An organization has an established change management program that includes standard operating procedures (SOPs). It wants to implement changes consistently and effectively.

What role do SOPs play in the change management process?

Answers:

***SOPs define routine operations or changes and provide detailed instructions for their implementation.**

SOPs help track and communicate the status and outcome of approved changes.

SOPs outline the steps for employees to follow when conducting an impact analysis.

SOPs provide guidelines for developing backout plans for changes.

Explanation:

Essential in change management, SOPs provide detailed instructions for routine operations and changes, ensuring consistency and effectiveness in implementation.

While SOPs may include guidelines for documenting and reporting changes, their primary role is not to track and communicate the status and outcome of approved changes.

SOPs do not cover impact analysis procedures separate from change management programs.

SOPs may guide developing backout plans, but it is not their primary role. Their primary focus is on defining the routine operations or changes and providing instructions for their implementation.

q_change_management_sop_02_secp8

What role do standard operating procedures (SOPs) play in the change management process of an organization's established change management program, which aims to guarantee consistent and effective implementation of changes?

Answers:

***SOPs outline regular operations or modifications and provide specific guidance for their execution.**

SOPs help track and communicate the status and outcome of approved changes.

SOPs outline the steps for employees to follow when conducting an impact analysis.

SOPs provide guidelines for developing backout plans for changes.

Explanation:

Having SOPs is crucial for change management since they offer precise guidelines for regular operations and changes. This guarantees that implementation is consistent and effective.

While SOPs may provide instructions for documenting and reporting changes, their primary purpose is not to monitor and inform about the progress and result of approved changes.

The SOPs do not include specific impact analysis procedures distinct from change management programs.

Even though SOPs may include instructions for creating backout plans, their primary purpose is to define routine operations and changes and provide guidance on implementing them.

q_change_management_stakeholders_secp8

Upon receiving the findings from a recent inspection, a senior technician must identify the various parties needed to implement change management solutions.

Based on those impacted by the change, who are primary stakeholders within the change management spectrum? (Select three.)

Answers:

***Partners**

***Vendors**

***Change Advisory Board (CAB)**

Project manager

Third-party suppliers

Customers

Consultants

Explanation:

Stakeholders in the change management process describe the individuals or groups impacted (or interested) in the change. Based on the scenario, the following are stakeholders within the change management spectrum:

The role of partners.

The vendors tied to the project.

The change advisory board (CAB) is also a stakeholder in the change management process and uses the information provided by the ownership cluster to make informed decisions.

Contrary to being a stakeholder, the following are roles that are not involved within the change management spectrum:

A project manager is part of the ownership cluster for change management and is primarily responsible for implementing specific changes.

Because primary stakeholders are normally within the company, third-party suppliers, customers, and consultants are not usually included as primary stakeholders in a change management solution.

q_change_management_version_control_01_secp8

Change management is not just for implementing software updates or hardware changes.

For example, version control refers to capturing changes made to important documents a company needs.

What are some documents that would utilize version control? (Select three.)

Answers:

***Code**

Faxes

***Diagrams**

***Important data**

Financial records

Transactional documents

Employee timecards

Explanation:

The following are documents that might utilize version control:

Capturing implemented changes in code is important as it will allow for a quick reversion to a known good state of the code if a change causes problems.

Capturing changes within diagrams ensures only the most recent diagrams will be available while archiving the previous documents, but still having them available for reference.

Important data will vary from company to company, but tracking changes will allow references for only the most recent data and avoid confusion.

The following are documents that would not normally fall under change management and version control:

Faxes and employee timecards do not fall under change management as they are one-time communications and do not need tracking or logging.

Financial records are meant to be stable and fixed reflections of financial transactions without changes (except for the occasional correction).

Transactional documents are also one-time communications that should reflect circumstances around a transaction at the time of the event.

q_change_management_version_control_02_secp8

How can a cybersecurity analyst effectively utilize version control to maintain a historical record of changes and ensure security in the organization's IT systems and applications?

Answers:

***Use version control to track changes in network diagrams and configuration files.**

Implement version control only for critical documents and code.

Revert to previous versions of documents without assessing their impact on security.

Use version control solely for policy updates, neglecting changes to code.

Explanation:

The analyst should use version control to track changes in critical documents and code, network diagrams, and configuration files. This comprehensive approach helps maintain a historical record of changes and ensures security across various aspects of the organization's IT infrastructure.

Implementing version control solely for critical documents and code would be insufficient to maintain a comprehensive historical record and secure the IT systems.

Reverting to previous versions of documents without assessing their impact on security may introduce vulnerabilities.

Using version control only for policy updates and neglecting changes to code does not address the importance of tracking all relevant areas.

11.3 Automation and Orchestration

As you study this section, answer the following questions:

What roles do automation and scripting play in IT operations, specifically security management?

How does an orchestrated system function if a threat is detected?

How does automation help employees?

What are five challenges that automation presents to maintaining security systems?

The key terms for this section include:

Term	Definition
Workforce multiplier	A tool or automation that increases employee productivity, enabling them to perform more tasks to the same standard per unit of time.
Reaction times	The elapsed time between an incident occurring and a response being implemented.
Single point of failure	A component or system that would cause a complete interruption of a service if it failed.

Technical debt	Costs accrued by keeping an ineffective system or product in place, rather than replacing it with a better-engineered one.
Standard configurations	In an IaC architecture, the property that an automation or orchestration action always produces the same result, regardless of the component's previous state.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.7 Explain the importance of automation and orchestration related to secure operations. Use cases of automation and scripting Ticket creation Enabling/disabling services and access Continuous integration and testing Benefits Efficiency/time saving Enforcing baselines Standard infrastructure configurations Staff retention Reaction time Workforce multiplier Other considerations Complexity Cost Single point of failure Technical debt Ongoing supportability

11.3.1 Automation and Scripting (Lesson Video)

Transcript:

Think about the chores you do every day, like washing your clothes or doing the dishes. You could do these things by hand if you wanted, but it would take longer and probably be less effective than doing them with a machine. The process of having machines assist human labor is called automation, and a similar mindset is used in many kinds of computing. In this lesson, I'm going to cover a variety of facts about automation so you can have a better idea of how this works.

The term Development, Security, and Operations is more commonly referred to as DevSecOps. This is the philosophy in which everyone in an organization is responsible for their system's security. This means that everyone should implement security decisions in the same way that development and operation decisions are made. If everyone in a company becomes better at security, the entire organization benefits, and customers have greater software assurance.
While this might take a while to set up initially, it eventually streamlines the process for everyone and leads to lower costs and development times. During the software development process, engineers constantly check each other's work and test the software for bugs and vulnerabilities. In this way, fewer vulnerabilities make their way into the final product, which means that we end up investing less time and money into fixing things post-release.
In a musical orchestra, there are many different instruments that are brought together to create something wonderful. Members follow what's written on their music sheets and the real-time directions given to them by the conductor. Just like a musical orchestra, workflows can be orchestrated to make them as efficient as possible so that they produce the best results.
A workflow refers to a collection of tasks that are performed in a logical sequence. Orchestration means that you plan these tasks in such a way that they're as efficient as possible.
For example, you can orchestrate different parts of your workflow, including the development, quality assurance, and security. Orchestration often incorporates the use of tools that automatically complete certain tasks in a sequence.
To help with workflow orchestration, many companies are using automated cybersecurity solutions that are able to quickly identify and resolve potential attacks. One of these solutions is Security Orchestration, Automation, and Response, which is known by the acronym SOAR.
SOAR refers to a collection of software programs that allow an organization's security team to collect various inputs they can monitor. The point of SOAR is to replace tasks that are repetitive and done manually with automated workflows. These systems automatically flag security incidents and respond to them in a predetermined way.
This means that these incidents are caught earlier. This also frees up security analysts to spend their time and attention on only the most advanced security threats. For example, using a firewall application, SOAR can automatically detect a brute-force login attack and block the attacker's IP address.
Another orchestration technique we see is known as continuous integration, or CI. When you work on a software development team, you often have multiple people contributing to the same project. The practice of continuous integration means that you automate all integration changes made by these contributors back into a shared mainline.
In general, there's a central repository where all code changes are merged into a single file that's used to test the current build's effectiveness. Under this strategy, developers try to merge their changes back to the main branch as often as possible. The new changes are then automatically tested to make sure they don't break the application when they're integrated. This can greatly streamline the development process because developers don't need to manually discuss the changes they make with the rest of the team. This also cuts down on overhead costs and lets developers focus on the code and not on a complex web of communication.
Continuous delivery is like an extension of continuous integration. It automatically deploys all the changes coders make into a production environment. You have both an automated testing process and an automated release process that you can set to occur at whichever interval you feel is best.
Continuous deployment goes a step further than continuous delivery. Continuous deployment means that any change that goes through all the production pipeline stages is automatically released to customers. There's no human intervention in this process, and only if the change fails one of the tests along the way is it prevented from being pushed out. This has the potential to create an extremely streamlined process and quicker responses to customer feedback.
That's it for this lesson. In this lesson, we first discussed DevSecOps and explained how this philosophy adds security measures to every step of the development process. Next, we discussed workflow orchestration and showed how it can help the development process run more smoothly with all the different elements aligned in the most efficient way. A big part of this is automating whichever elements can possibly be automated.

We also introduced you to some tools to help in your orchestration, including SOARS, which lets the security team automatically collect data to help them identify security threats. And finally, we discussed how continuous integration lets you automate the integration of changes from multiple developers into a central staging area. Lastly, we showed that you can use continuous deployment to ensure changes are automatically deployed to a production environment and pushed out to end users.

11.3.2 Automation and Scripting Facts

Automation and orchestration are powerful tools for managing security operations. Automation uses software to perform repetitive, rule-based tasks, such as monitoring for threats, applying patches, maintaining baselines, or responding to incidents, to improve efficiency and reduce the likelihood of human error. Orchestration enhances automation by coordinating and streamlining the interactions between automated processes and systems. Orchestration supports seamless and integrated workflows, especially in large, complex environments with many different security tools and systems. Automation and orchestration also provide clear audit trails supporting regulatory compliance and incident investigation. While their implementation comes with challenges such as complexity, cost, and the potential for a single point of failure, careful management of these tools can greatly improve an organization's security posture.

This lesson covers the following topics:

- Automation and scripting
- Benefits of automation and orchestration in security operations
- Important considerations
- Benefits of infrastructure management automation

Automation and Scripting

Automation and scripting have emerged as critical tools in modern IT operations, helping organizations streamline processes, enhance security, and improve efficiency. Automation serves as a tool to enhance both security governance and change management. In terms of governance, automation can help enforce security policies more consistently and efficiently, and it can aid in monitoring and reporting to provide valuable insights for leadership teams and risk managers. In change management, automation can reduce the risk of human error, reduce implementation time, and provide clear audit trails. For example, scripts are effective for applying patches and updates across an organization's systems uniformly, and automation tools can track these changes for later review.

Benefits of Automation and Orchestration in Security Operations

Automation and orchestration also offer many important benefits to security operations. Primarily, they enhance efficiency by enabling repetitive tasks to be performed quickly and consistently, reducing the burden on security teams and minimizing the likelihood of human error (sometimes referred to as a workforce multiplier).

Operator fatigue refers to the mental exhaustion experienced by cybersecurity professionals due to their work's continuous, high-intensity nature. Security analysts must monitor numerous systems for potential threats, manage high volumes of alerts (including many false positives,) and respond to confirmed threats as quickly as possible. These working conditions often lead to long hours, anxiety, and elevated stress levels, resulting in operator fatigue. This fatigue is a significant concern in cybersecurity because it can lead to decreased alertness and cognitive function and impair the ability of security personnel to identify and respond to threats effectively. Fatigue results in missed critical alerts, slower response times, and a greater likelihood of errors, any of which can compromise security.

Automation and orchestration play crucial roles in combating operator fatigue in security operations by minimizing the repetitive, manual tasks that often contribute to operator fatigue. Automation and orchestration significantly reduce a security team's

workload by automating routine tasks, such as scanning for vulnerabilities, applying patches, or monitoring systems for anomalous activities. This allows for the more efficient use of resources and frees security personnel to focus on more complex, strategic issues that require human judgment and creativity rather than repetitive tasks. Orchestration enhances the impact of automation by coordinating automated tasks across different systems and software tools and reduces detection and reaction times .

For example, if a threat is detected, an orchestrated system can automatically isolate the affected subnet, perform basic analysis and reporting, notify security teams, generate tickets, and document the incident, all without human intervention. Other benefits of automation include enforcing standardized baselines through configuration management tools to override unauthorized changes made to endpoints automatically. A standard baseline in configuration management is a well-defined set of approved configurations and settings that serve as a reference point for establishing and maintaining the desired state of a system. Automation and orchestration can significantly alleviate operator fatigue by reducing the volume of manual, routine tasks and improving the efficiency of security operations, leading to greater job satisfaction, increased alertness and effectiveness in threat detection and response, and, ultimately, more robust security operations.

Automation can support staff retention initiatives by reducing fatigue from repetitive tasks. Automation practices can free staff to perform more rewarding work and increase job satisfaction.

Important Considerations

While automation and orchestration provide numerous benefits, they also present some significant challenges, some of which are listed below:

Complexity — Implementing automation and orchestration requires a deep understanding of an organization's systems, processes, and interdependencies. A poorly planned or executed automation strategy can add complexity, making systems more difficult to manage and maintain.

Cost — The initial cost of implementing automation and orchestration can be high, including costs associated with acquiring and developing appropriate tools, integrating them into existing systems, and training staff to use them effectively. Automation software maintenance and upgrades can also be costly.

Single Point of Failure — If a critical automated system or process fails, it could impact multiple areas of the organization, causing widespread problems.

Technical Debt — Organizations can accrue technical debt if automation and orchestration tools are implemented hastily, resulting in poorly documented code, "brittle" system integrations, or poor maintenance. Over time, this debt can lead to system instability, complexity, and increased costs, ironically similar to the problems associated mainly with legacy systems.

Ongoing Support — Automation and orchestration systems require ongoing support to stay effective and secure, including updates and patches, reviewing and improving automated processes, and continuous education. Without adequate support, the benefits of automation and orchestration are quickly eroded.

Maintaining system security when new hardware or infrastructure items are added to the network can be achieved by enforcing standard configurations across the company. With automated configurations, these newly added items can be kept up-to-date and secure.

Benefits of Infrastructure Management Automation

Automating and orchestrating infrastructure configurations introduces numerous benefits. Enforcing standardized configurations ensures consistency and accuracy throughout the infrastructure. Automation saves time and resources by allowing configurations to be quickly deployed, and it also enhances scalability and flexibility by simplifying the deployment and configuration of new resources.

Furthermore, automation and orchestration improve standardization, compliance, and change management by enforcing predefined configuration standards, making auditing and change tracking easier, and controlling configuration drift. Additionally, automation can strengthen security and governance by enforcing security controls, applying patches consistently, and automating security-related tasks.

11.3.3 Practice Questions (Section Quiz)

q_auto_scripting_automation_challenges_secp8

A tech director evaluates the benefits of implementing automation and orchestration into the organization after receiving approval and funding notification for the annual budget.

Knowing several benefits tied to automation, what challenges exist when managing automation? (Select three.)

Answers:

***Technical debt**

***Single point of failure**

***Cost**

Staff retention

Risk of human error

Implementation time

Applying patches and updates

Explanation:

The following are automation challenges:

One challenge is that organizations can accrue technical debt if the implementation of automation and orchestration tools begins in haste, resulting in poorly documented code, "brittle" system integrations, or poor maintenance.

A challenge can include challenges associated with single points of failure, such as a critical automated system or process failure, impacting multiple areas of the organization and causing widespread problems.

A challenge can involve funding, recognizing that the initial cost of implementing automation and orchestration can be high, including costs associated with acquiring and developing appropriate tools and integrating them into existing systems.

The following are not normally associated with automation challenges:

While important to an organization, they do not consider staff retention a challenge to automation but rather a tool to reduce employee/operator fatigue.

Benefits of using automation include reducing the risk of human error, reducing implementation time, and is effective for applying patches and updates.

q_auto_scripting_automation_orchestration_benefit_01_secp8

As a Security Operations Center (SOC) analyst for a large financial institution that deals with high volumes of alerts and potential threats, what crucial benefit does implementing automation and orchestration in security operations provide?

Answers:

Automation and orchestration simplify the nature of threats and reduce the volume of alerts.

***Automation and orchestration enable repetitive tasks to be performed quickly and consistently, minimizing human error.**

Automation and orchestration eliminate the need for human intervention in security operations.

Automation and orchestration help to cut costs by reducing the number of cybersecurity professionals needed.

Explanation:

Automation improves efficiency in security operations by quickly and consistently performing repetitive tasks, which reduces the incidence of human error often associated with such tasks.

Automation and orchestration do not simplify the nature of threats nor reduce the volume of alerts. They provide tools to handle them more efficiently and reduce manual workload.

Human intervention remains a key part of security operations. Automation and orchestration merely augment human roles, freeing the team to focus on complex issues requiring human judgment.

While automation can reduce some workloads, it does not necessarily decrease the need for cybersecurity professionals. Expert human judgment is essential for streamlining and verifying automated systems.

q_auto_scripting_automation_orchestration_benefit_02_secp8

A technical consultant reviews available automation and orchestration options for security operations after realizing employees' actions place the network architecture at risk.

What are the benefits associated with automation and orchestration implementation? (Select three.)

Answers:

***It enhances efficiency by enabling repetitive tasks to perform with mitigation of risk to human error consistently.**

***It coordinates automated tasks across different systems and software tools for quicker response.**

***It assists staff members from experiencing fatigue and enhances opportunities for retention.**

It assists in funding needed orchestration and automation at a higher price point.

It requires a deep understanding of an organization's systems, processes, and interdependencies.

Ongoing support to stay effective and secure.

A single point of failure for the system.

Explanation:

The following are benefits associated with automation and orchestration implementation:

Automation and orchestration offer many benefits to security operations. They enhance efficiency by enabling users to quickly and consistently perform repetitive tasks, reducing the burden on security teams and minimizing the likelihood of human error.

Orchestration enhances the impact of automation by coordinating automated tasks across different systems and software tools, reducing detection and reaction times.

Automation supports staff retention initiatives by reducing fatigue from repetitive tasks. Automation practices can free staff to perform more rewarding work and increase job satisfaction.

The following are significant challenges when using an automation and orchestration implementation:

Cost is not a benefit but more of a challenge as the initial cost of implementing automation and orchestration can be high, including costs associated with acquiring and developing appropriate tools and training staff to use them effectively.

Implementing automation and orchestration requires a deep understanding of an organization's systems, processes, and interdependencies. A poorly planned or executed automation strategy can add complexity, making systems more difficult to manage and maintain.

Automation and orchestration systems require ongoing support to stay effective and secure, including updates and patches, reviewing and improving automated processes, and continuous education. Without adequate support, the benefits of automation and orchestration are quickly eroded.

A single point of failure means that if a critical automated system or process fails, it could impact multiple areas of the organization, causing widespread problems.

q_auto_scripting_automation_orchestration_challenge_secp8

Upon receiving additional funding for the new quarter, a software team leader looks to acquire new automation and orchestration tools to enhance the IT department.

What is NOT considered a benefit of automation and orchestration implementation for infrastructure management?

Answers:

***Enforcing standardized baselines through configuration management tools**

Enforcing standardized configurations to ensure consistency

Saving time and resources by allowing configurations to deploy quickly

Enhancing scalability and flexibility by simplifying deployment

Explanation:

A benefit of automation in security operations, and not infrastructure management, is to enforce standardized baselines through configuration management tools. It overrides unauthorized endpoint changes automatically.

Automating and orchestrating infrastructure configurations introduces numerous benefits. Enforcing standardized configurations ensures consistency and accuracy throughout the infrastructure.

Automation in infrastructure management saves time and resources by allowing configurations to deploy quickly.

Automation in infrastructure management enhances scalability and flexibility by simplifying the deployment and configuration of new resources. Furthermore, it improves compliance and makes auditing and change tracking easier, which controls configuration drift.

q_auto_scripting_automation_scripting_secp8

A multinational corporation is upgrading its IT infrastructure to enhance security governance and streamline its change management process. The IT department is considering various strategies to accomplish this update.

Which strategy MOST effectively achieves the corporation's goals, considering the inherent risks and benefits?

Answers:

***Implementing automation and scripting to perform tasks quickly and efficiently**

Manual monitoring of security controls and change management protocols

Outsourcing the entire IT operations to a third-party vendor

Using proprietary security solutions without automation

Explanation:

Automation and scripting enhance efficiency and security, which aligns with the corporation's goals. The IT department should use automation and scripting with care to prevent potential risks.

Automation improves monitoring efficiency by making it faster, better, and cheaper than manual monitoring. Manual monitoring does not help the corporation achieve its goal of making its processes more efficient. Instead, it slows down the process, lowers the quality, and increases the cost.

Outsourcing may fail to directly address the corporation's security governance and change management needs. It may also lead to complexities and loss of control.

Proprietary solutions without automation lack efficiency. They do not align with the corporation's objectives for streamlined processes.

q_auto_scripting_automation_tools_secp8

A financial institution is evaluating its incident response plan and wants to incorporate automation to accelerate the detection and mitigation of security breaches.

The security team must ensure that the automation does not inadvertently cause additional issues or conflicts.

What is the BEST approach the team should employ when incorporating automation?

Answers:

Deploying automation scripts without testing or validation.

Utilizing a manual-only approach without integrating any automation tools.

***Integrating automation tools with real-time monitoring and alerting capabilities.**

Outsourcing automation development to a non-specialized third-party vendor.

Explanation:

Integrating automation tools with real-time monitoring and alerting provides timely detection and mitigation of security breaches. This approach aligns with the financial institution's goal of accelerating incident response without compromising security.

Deploying scripts without testing or validation may lead to conflicts or additional issues, failing to meet the institution's requirements for secure automation.

Utilizing a manual-only approach without integrating any automation tools is not an automation solution.

Outsourcing automation development to a non-specialized third-party vendor poses a risk of introducing additional issues and conflicts.

q_auto_scripting_code_branch_secp8

A tech department evaluates the benefits of automation and scripting after recently acquiring new funding.

What capability within automation and scripting allows developers to regularly merge their changes back to the main code branch and evaluate each merge automatically to help detect and fix integration problems?

Answers:

***Continuous integration and testing**

Guardrails

User provisioning

Resource provisioning

Explanation:

The principles of continuous integration and testing hinge heavily on automation. In this approach, developers regularly merge their changes back to the main code branch and evaluate each merge automatically to help detect and even fix integration problems.

Guardrails and security groups are not the optimal answer in this scenario. Instead, they provide frameworks for managing security within an organization.

User provisioning describes creating, modifying, or deleting user accounts and access rights across IT systems.

Resource provisioning describes allocating IT resources such as servers, storage, and networks to applications and users.

q_auto_scripting_complexity_secp8

A technician wants to implement automation within the team's workspace.

How does complexity impact automation and orchestration?

Answers:

***Poorly planned strategies can make systems difficult to maintain.**

It can impact multiple areas of the organization, causing widespread problems.

It can result in poorly documented code, leading to instability and increased costs.

It can quickly erode if they do not continue the needed patches and updates.

Explanation:

While automation and orchestration provide numerous benefits, they can also present numerous challenges. A poorly planned or executed automation strategy can add complexity, making systems difficult to maintain.

Single points of failure can impact multiple areas of the organization, causing widespread problems since no redundancy of operations exists.

When organizations make haste decisions in technology, technical debt can result in poorly documented code or maintenance. Over time, technical debt can lead to system instability and increased costs.

Automation and orchestration require ongoing support to stay effective and secure. Automation and orchestration will quickly erode without support.

q_auto_scripting_infrastructure_management_secp8

Which of the following is a benefit of infrastructure management automation?

Answers:

Increasing the volume of manual, routine tasks

Decreasing the efficiency of security operations

***Reducing the risk of human error**

Increasing operator fatigue

Explanation:

Reducing the risk of human error is a benefit of infrastructure management automation. Automation can help reduce the risk of human error by automating repetitive tasks, such as applying patches and updates across an organization's systems uniformly.

Increasing the volume of manual, routine tasks is not a benefit of infrastructure management automation. In fact, one of the main advantages of automation is that it reduces the volume of manual, routine tasks, allowing staff to focus on more complex, strategic issues.

Decreasing the efficiency of security operations is not a benefit of infrastructure management automation. Automation and orchestration enhance efficiency by enabling repetitive tasks to be performed quickly and consistently, reducing the burden on security teams and minimizing the likelihood of human error.

Increasing operator fatigue is not a benefit of infrastructure management automation. By reducing the volume of manual, routine tasks, automation can significantly alleviate operator fatigue, leading to greater job satisfaction and increased alertness and effectiveness in threat detection and response.

q_auto_scripting_operator_fatigue_secp8

A company has been experiencing issues with operator fatigue within the cybersecurity team, leading to decreased alertness and cognitive function.

Considering different strategies to help combat this issue, how can automation and orchestration assist in addressing operator fatigue in security operations?

Answers:

By reducing the initial cost of implementing security measures.

By enforcing standardized baselines and overriding unauthorized changes.

***By automating routine tasks, allowing cybersecurity personnel to focus on more complex, strategic issues.**

By increasing the complexity of the company's systems and processes.

Explanation:

By automating routine tasks such as scanning for vulnerabilities, applying patches, or monitoring systems, automation and orchestration can significantly reduce a cybersecurity team's workload. This reduces operator fatigue by allowing cybersecurity personnel to focus on more complex, strategic issues.

While automation and orchestration might save costs in the long run due to increased efficiency and reduced errors, the initial cost of implementation can be high.

While enforcing standardized baselines and overriding unauthorized changes is a benefit of automation and orchestration, it does not directly address operator fatigue. This benefit maintains consistent system configurations and security.

Increased complexity is actually a potential challenge of implementing automation and orchestration, not a benefit, and does not address operator fatigue.

q_auto_scripting_security_analysts_secp8

In an IT environment, automation and scripting play a critical role in managing services and access.

How does automation assist security analysts in their daily tasks?

Answers:

By helping in user and resource provisioning.

By improving the efficiency of ticketing platforms.

By facilitating the development of more complex systems such as SOAR platforms.

***By enabling and disabling services, modifying access rights, and maintaining the lifecycle of IT resources.**

Explanation:

Automation and scripting are essential tools for managing services and access within an IT environment. This includes enabling or disabling services, modifying access rights, and maintaining the lifecycle of IT resources, which directly aligns with the tasks of security analysts.

While automation assists in user and resource provisioning, the question specifically asks how automation assists security analysts in managing services and access, not provisioning.

While improving the efficiency of ticketing platforms is a benefit of automation, it does not directly apply to the tasks of security analysts in managing services and access within an IT environment.

While automation does facilitate the development of more complex systems like Security Orchestration Automated Response (SOAR) platforms, this is not a direct way that it assists security analysts in managing services and access.

q_auto_scripting_technical_debt_secp8

A cybersecurity analyst works in an organization with several legacy systems with undocumented code and poorly maintained integrations.

How can the cybersecurity analyst address the technical debt associated with these legacy systems using automation and orchestration?

Answers:

***Utilize automation and orchestration to improve documentation and maintenance of the code and integrations.**

Avoid implementing automation and orchestration to prevent system instability.

Apply automated patches and updates without understanding the systems' complexities.

Ignore technical debt as it poses no immediate security risk.

Explanation:

The cybersecurity analyst can streamline and standardize the documentation process for the legacy system's code and integrations by implementing automation and orchestration. Automated processes can generate comprehensive documentation, making it easier for the analyst and other team members to understand and modify the code when necessary.

Avoiding implementation is incorrect because implementing automation and orchestration focusing on documentation and maintenance can reduce technical debt, leading to better system stability.

Applying automated patches is incorrect because applying automated patches and updates without understanding the systems' complexities may exacerbate technical debt and introduce potential security risks.

Ignoring technical debt can lead to long-term problems, including system instability and increased security risks.

q_auto_scripting_ticketing_advantages_secp8

What advantages can automation and scripting bring to IT operations ticketing platforms?

Answers:

***Support tickets are automatically generated and routed for incidents detected by monitoring systems.**

Security policies are enforced to prevent risky activities and unauthorized behavior.

Services and access are better managed within an IT environment.

Automated testing helps improve code quality and speeds up development cycles.

Explanation:

Automation and scripting can create support tickets for IT incidents, ensuring that the relevant team or individual can promptly resolve them.

In IT operations, automation is a valid means of enforcing security policies to prevent risky activities and unauthorized behavior, although it is not limited to ticketing platforms.

Although not directly tied to ticketing platforms, managing services and access within an IT environment is a valuable application of automation.

Enhancing code quality and expediting development cycles through automatic testing are valuable benefits of automation, but they are not directly associated with ticketing platforms in IT operations.

q_auto_scripting_workforce_multiplier_secp8

A third-party escalation team participates in a newly contracted project with numerous cyber teams. Being unfamiliar with cyberspace, the escalation team struggles to understand concepts and naming conventions.

What is automation and orchestration also known as?

Answers:

***Workforce multiplier**

Guardrail

User provisioning

Resource provisioning

Explanation:

Automation and orchestration, also known as a workforce multiplier, enhances efficiency by quickly and consistently performing on enabling repetitive tasks, reduces the burden on security teams, and minimizes the likelihood of human error.

Guardrails and security groups are not the optimal answer in this scenario, but they provide frameworks for managing security within an organization.

User provisioning describes creating, modifying, or deleting user accounts and access rights across IT systems and is not a reference to automation and orchestration.

Resource provisioning describes allocating IT resources such as servers, storage, and networks to applications and users.

12.0 Risk Management Processes

12.1 Risk Management Processes and Concepts

As you study this section, answer the following questions:

Why are disaster recovery policies important for an organization's security?

What is the difference in acceptance and mitigation in risk management?

What is the difference in qualitative and quantitative risk assessment?

How is the annualized rate of occurrence (ARO) calculated?

What are examples of external risk types?

The key terms for this section include:

Term	Definition
Risk management	The cyclical process of identifying, assessing, analyzing, and responding to risks.
Business impact analysis (BIA)	Systematic activity that identifies organizational risks and determines their effect on ongoing mission-critical operations.
Mission essential function (MEF)	Business or organizational activity that is too critical to be deferred for anything more than a few hours, if at all.
Maximum tolerable downtime (MTD)	The longest period that a process can be inoperable without causing irrevocable business failure.
Recovery time objective (RTO)	The maximum time allowed to restore a system after a failure event.
Work recovery time (WRT)	In disaster recovery, time additional to the RTO of individual systems to perform reintegration and testing of a restored or upgraded system following an event.
Recovery point objective (RPO)	The longest period that an organization can tolerate lost data being unrecoverable.
Mean time between failures (MTBF)	Metric for a device or component that predicts the expected time between failures.

Mean time to repair (MTTR)	Metric representing average time taken for a device or component to be repaired, replaced, or otherwise recover from a failure.
Risk identification	Within overall risk assessment, specific process of listing sources of risk due to threats and vulnerabilities.
Risk mitigation (or remediation)	The response of reducing risk to fit within an organization's willingness to accept risk.
Risk deterrence (or reduction)	In risk mitigation, the response of deploying security controls to reduce the likelihood and/or impact of a threat scenario.
Avoidance	In risk mitigation, the practice of ceasing activity that presents risk.
Risk acceptance	The response of determining that a risk is within the organization's appetite and no countermeasures other than ongoing monitoring is needed.
Risk exception	A category of risk management that uses alternate mitigating controls to control an accepted risk factor.
Risk exemption	A category of risk management that accepts an unmitigated risk factor.
Residual risk	Risk that remains even after controls are put into place.
Likelihood	In risk calculation, the chance of a threat being realized, expressed as a percentage.
Probability	The mathematical measure of the possibility of a risk occurring.
Impact	The severity of the risk if realized by factors such as the scope, value of the asset, or the financial impacts of the event.
Enterprise risk management (ERM)	The comprehensive process of evaluating, measuring, and mitigating the many risks that pervade an organization.
Risk assessment	The process of identifying risks, analyzing them, developing a response strategy for them, and mitigating their future impact.
Risk analysis	Process for qualifying or quantifying the likelihood and impact of a factor.

Quantitative risk analysis	A numerical method that is used to assess the probability and impact of risk and measure the impact.
Single loss expectancy (SLE)	The amount that would be lost in a single occurrence of a particular risk factor.
Annualized loss expectancy (ALE)	The total cost of a risk to an organization on an annual basis. This is determined by multiplying the SLE by the annual rate of occurrence (ARO).
Annualized rate of occurrence (ARO)	In risk calculation, an expression of the probability/likelihood of a risk as the number of times per year a particular loss is expected to occur.
Qualitative risk analysis	The process of determining the probability of occurrence and the impact of identified risks by using logical reasoning when numeric data is not readily available.
Inherent risk	Risk that an event will pose if no controls are put in place to mitigate it.
Risk register	A document highlighting the results of risk assessments in an easily comprehensible format (such as a 'traffic light' grid). Its purpose is for department managers and technicians to understand risks associated with the workflows that they manage.
Heat map risk matrix	A graphical table indicating the likelihood and impact of risk factors identified for a workflow, project, or department for reference by stakeholders.
Key Risk Indicators (KRIs)	The method by which emerging risks are identified and analyzed so that changes can be adopted to proactively avoid issues from occuring.
Risk owner	An individual who is accountable for developing and implementing a risk response strategy for a risk documented in a risk register.
Risk appetite	A strategic assessment of what level of residual risk is tolerable for an organization.
Risk tolerance	Determines the thresholds that separate different levels of risk.
Risk reporting	A periodic summary of relevant information about a project's current risks. It provides a summarized overview of known risks, realized risks, and their impact on the organization.
Continuity of operations (COOP)	Identifies how business processes should deal with both minor and disaster-level disruption by ensuring that there is processing redundancy supporting the workflow.

Capacity planning	A practice which involves estimating the personnel, storage, computer hardware, software, and connection infrastructure resources required over some future period of time.
Hot site	A fully configured alternate processing site that can be brought online either instantly or very quickly after a disaster.
Warm site	An alternate processing location that is dormant or performs noncritical functions under normal conditions, but which can be rapidly converted to a key operations site if needed.
Cold site	A predetermined alternate location where a network can be rebuilt after a disaster.
Geographic dispersion	A resiliency mechanism where processing and data storage resources are replicated between physically distant sites.
Platform diversity	Cybersecurity resilience strategy that increases attack costs by provisioning multiple types of controls, technologies, vendors, and crypto implementations.
Tabletop exercises	A discussion of simulated emergency situations and security incidents.
Simulations	A testing technique that replicates the conditions of a real-world disaster scenario or security incident.
Parallel processing tests	Running primary and backup systems simultaneously to validate the functionality and performance of backup systems without disrupting normal operations.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	3.4 Explain the importance of resilience and recovery in security architecture. High availability Site considerations Hot Cold Warm

- Geographic dispersion
- Platform diversity
- Multi-cloud systems
- Continuity of operations
- Capacity planning
  - People
  - Technology
  - Infrastructure
- Testing
  - Tabletop exercises
  - Fail over
  - Simulation
  - Parallel processing
- Backups

4.2 Explain the security implications of proper hardware, software, and data asset management.

- Assignment/accounting

5.1 Summarize elements of effective security governance.

- Policies
  - Business continuity
  - Incident response

5.2 Explain elements of the risk management process.

- Risk identification
- Risk assessment
  - Ad hoc

- Recurring
- One-time
- Continuous

Risk analysis

- Qualitative
- Quantitative
- Single loss expectancy (SLE)
- Annualized loss expectancy (ALE)
- Annualized rate of occurrence (ARO)
- Probability
- Likelihood
- Exposure factor
- Impact

Risk register

- Key risk indicators
- Risk owners
- Risk threshold

Risk tolerance

Risk appetite

- Expansionary
- Conservative
- Neutral

Risk management strategies

- Transfer
- Accept

	Exemption Exception Avoid Mitigate Risk reporting Business impact analysis Recovery time objective (RTO) Recovery point objective (RPO) Mean time to repair (MTTR) Mean time between failures (MTBF)

12.1.1 Risk Types and Tolerance (Lesson Video)

Transcript:

Risk management is the process of identifying vulnerabilities and threats, and then deciding on countermeasures to reduce the risk to an acceptable level. In other words, risk management is when a person or group considers worst-case scenarios and determines just how bad those scenarios would be for their organization. Senior management is ultimately responsible for residual risk, which is what remains after taking risk-reduction measures.

This all might sound kind of depressing. But it's actually pretty liberating because you can feel secure knowing that you've thought through all possible problems. It's kind of like insurance. No one wants bad stuff to happen, but you invest in an insurance policy so that you can have peace of mind that if something does happen, you'll be ready.
The first step in risk management is to identify your organization's assets and how much each asset is worth. You want to pay special attention to assets that are critical for your business to function. You should place values on all assets, including the cost of lost data, failed systems, downtime, and training new employees.
Assets include physical items like computers and storage devices. These items are usually purchased, and so it's fairly easy to place a value on them. But it becomes a bit harder to place a value on an intangible asset like computer data or a seasoned employee's knowledge. While a computer can be easily replaced, the data itself is possibly far more valuable in the end. It's difficult to assign values to things like that, but it's important to attempt to.
After you've identified your assets, you should then identify possible asset threats, which can be internal or external. Internal threats come from within your organization and could include employee fraud, theft, system failure, sabotage, espionage, collusion, or snooping.
External threats are things like fire, water damage, burglars, internet attacks, market competition, or natural disasters.
Let's talk about risk analysis. Risk analysis is the practice of assessing which risks you've identified to be the most relevant to the organization. This is a key part of the risk management process. Risk assessment determines quantitative and qualitative risk values as they relate to a particular threat. As a best practice, we try to create risk assessments using quantitative measurements, which require us to assign a number value to each identified risk.
You should also evaluate your risks. This means determining if and when action should be taken in a worst-cast scenario. You would use this during the analysis phase to determine each risk's tolerability.
Finally, once all risks have been identified and assessed, their management falls into one of four categories: avoidance, transference, mitigation, and acceptance. Avoidance is the decision to avoid the risk altogether. In other words, you choose not to engage in an activity because the risk is too high.

Transference involves transferring the risk to someone else. Many organizations have opted to outsource their networking, security, and storage solutions. Doing so helps to transfer a bit of the risk to companies that are better prepared for potential threats. As businesses dependence on technology continues to grow, so does their cyber risk. Another way businesses help prepare for a potentially catastrophic loss is to take out cybersecurity insurance. As with most insurance policies, this protects against events that would have a big financial impact should they occur.
Mitigation is also an option. You could reduce a potential threat's risk by deploying security controls or other protections. An administrative control is an actionable procedure that should be followed to reduce risk. A technical control is a device that's used for this very purpose. For example, you could configure system redundancy so that if the original system were to go down, the redundant system could take over and continue providing services. Organizations should perform regular risk-control assessments to ensure that these controls continue to be effective.
Sometimes, organizations opt to accept a risk because the associated cost is acceptable, or the cost to protect the asset is just too high. A risk appetite statement, put simply, is the amount of risk an organization is willing to tolerate in order to reach its objectives. If you do decide to accept a risk, it's important to build response policies around a plausible outcome. Know that in some instances, there might actually be regulations that outline how certain risks must be lawfully handled. In these cases, you want to show that you took all the necessary steps to safeguard against loss. So make sure to study up and consider the guidance of legal counsel.
That's it for this lesson. In summary, the risk management process is the process for identifying threats and deciding which countermeasures to take to reduce an undesirable outcome. Asset and risk identification as well as risk analysis, evaluation, and management strategies all make up the risk management process.

12.1.2 Risk Types and Tolerance Facts

Risk management involves identifying potential issues, assessing their potential impact on the organization, and implementing controls to mitigate them. Key concepts include risk identification, risk assessment, mitigation, and monitoring. Risk appetite and risk tolerance are important in defining how much risk an organization is willing to accept. Methods such as ad hoc, recurring, one-time, or continuous risk assessments help organizations accurately understand their risks. Ultimately, effective risk management helps safeguard the organization's information assets, maintain regulatory compliance, and support strategic objectives.

This lesson covers the following topics:

- Risk Management Processes
- Identification of critical systems
- Mission essential functions
- Risk identification and assessment
- Risk management strategies
- Risk transference
- Risk acceptance

Risk Management Processes

Risk management is a process for identifying, assessing, and mitigating vulnerabilities and threats to the essential functions that a business must perform to serve its customers. You can think of this process as being performed over five phases:

- **Identify Mission Essential Functions** — mitigating risk can involve a large amount of expenditure so it is important to focus efforts. Effective risk management must focus on mission essential functions that could cause the whole business to fail if they are not performed. Part of this process involves identifying critical systems and assets that support these functions.

Identify Vulnerabilities — for each function or workflow (starting with the most critical), analyze systems and assets to discover and list any vulnerabilities or weaknesses to which they may be susceptible.

Identify Threats — for each function or workflow, identify the threat sources and actors that may take advantage of or exploit or accidentally trigger vulnerabilities.

Analyze Business Impacts — the likelihood of a vulnerability being activated as a security incident by a threat and the impact of that incident on critical systems are the factors used to assess risk. There are quantitative and qualitative methods of analyzing impacts and likelihood.

Identify Risk Response — for each risk, identify possible countermeasures and assess the cost of deploying additional security controls. Most risks require some sort of mitigation, but other types of response might be more appropriate for certain types and levels of risks.

Identification of Critical Systems

To support the resiliency of mission essential and primary business functions, it is crucial to perform an identification of critical systems. This means compiling an inventory of business processes and the assets that support them. Asset types include the following:

People (employees, visitors, and suppliers).

Tangible assets (buildings, furniture, equipment and machinery (plant), ICT equipment, electronic data files, and paper documents).

Intangible assets (ideas, commercial reputation, brand, and so on).

Procedures (supply chains, critical procedures, standard operating procedures).

For mission essential functions, it is important to reduce the number of dependencies between components. Dependencies are identified by performing a business process analysis (BPA) for each function. The BPA should identify the following factors:

Inputs — the sources of information for performing the function (including the impact if these are delayed or out of sequence).

Hardware — the particular server or datacenter that performs the processing.

Staff and other resources supporting the function.

Outputs — the data or resources produced by the function.

Process Flow — a step-by-step description of how the function is performed.

Business impact analysis (BIA) is a process that helps businesses understand the potential effects of disruptions on their operations. It involves identifying and assessing the impact of various unplanned threat scenarios on the business, such as accidents, emergencies, and disasters. By conducting a BIA, businesses can proactively create recovery strategies to minimize the impact of disruptions and ensure operational resilience.

For instance, if a DDoS attack suspends an e-commerce portal for five hours, the business impact analysis will be able to quantify the losses from orders not made and customers moving permanently to other suppliers based on historic data. The likelihood of a DoS attack can be assessed on an annualized basis to determine annualized impact in terms of costs. This informatioin is used to assess whether a security control, such as load balancing or managed DDoS mitigation, is worth the investment.

Mission Essential Functions

A mission essential function (MEF) is one that cannot be deferred. This means that the organization must be able to perform the function as close to continually as possible, and if there is any service disruption, the mission essential functions must be restored first.

Functions that act as support for the business or an MEF, but are not critical in themselves, are referred to as primary business functions (PBF).

Analysis of mission essential functions is generally governed by four main metrics:

Maximum tolerable downtime (MTD) is the longest period of time that a business function outage may occur for without causing irrecoverable business failure. Each business process can have its own MTD, such as a range of minutes to hours for critical functions, 24 hours for urgent functions, seven days for normal functions, and so on. MTDs vary by company and event. Each function may be supported by multiple systems and assets. The MTD sets the upper limit on the amount of recovery time that system and asset owners have to resume operations. For example, an organization specializing in medical equipment may be able to exist without incoming manufacturing supplies for three months because it has stockpiled a sizable inventory. After three months, the organization will not have sufficient supplies and may not be able to manufacture additional products, therefore leading to failure. In this case, the MTD is three months.

Recovery time objective (RTO) is the period following a disaster that an individual IT system may remain offline. This represents the amount of time it takes to identify that there is a problem and then perform recovery (restore from backup or switch to an alternative system, for instance).

Work Recovery Time (WRT) . Following systems recovery, there may be additional work to reintegrate different systems, test overall functionality, and brief system users on any changes or different working practices so that the business function is again fully supported.

RTO + WRT must not exceed MTD!

Recovery point objective (RPO) is the amount of data loss that a system can sustain, measured in time. That is, if a database is destroyed by a virus, an RPO of 24 hours means that the data can be recovered (from a backup copy) to a point not more than 24 hours before the database was infected. RPO is determined by identifying the maximum acceptable data loss an organization can tolerate in the event of a disaster or system failure and is established by considering factors such as business requirements, data criticality, and regulatory or contractual obligations. The calculation of RPO directly impacts the frequency of data backups, data replication requirements, recovery site selection, and technologies that support failover and high availability.

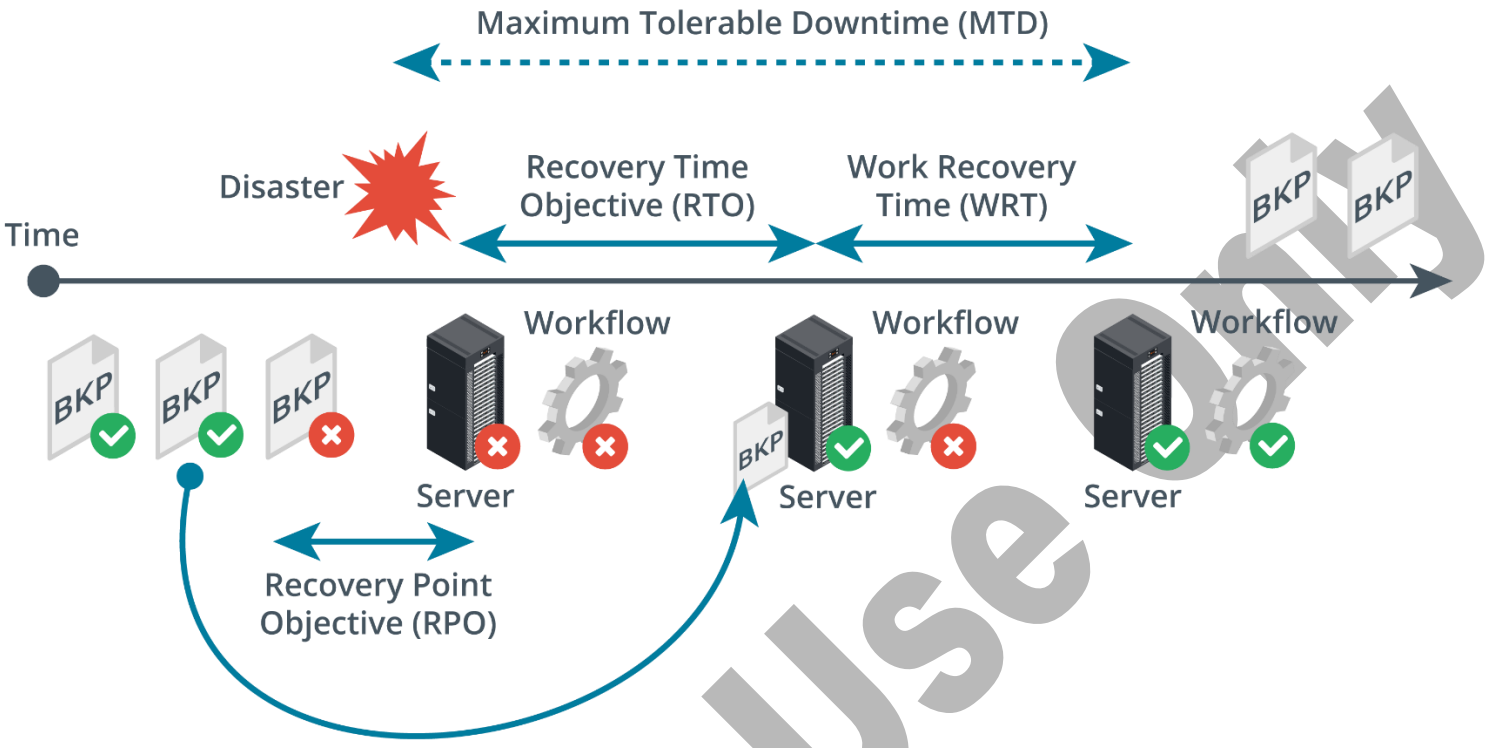

Metrics governing mission essential functions. (Images © 123RF.com.)

For example, a customer relationship management database might be able to sustain the loss of a few hours' or days' worth of data because employees can generally remember who they have contacted and the conversations they had over this time span. Conversely, order processing is generally more time sensitive, as data losses will represent lost orders, and it may be impossible to recapture them or the related processes initiated by order processing systems, such as accounting and fulfillment data.

MTD and RPO help to determine which business functions are critical and also to specify appropriate risk countermeasures. For example, if your RPO is measured in days, then a simple tape backup system should suffice; if RPO is zero or measured in minutes or seconds, a more expensive server cluster backup and redundancy solution will be required.

Mean time to repair (MTTR) and mean time between failures (MTBF) are key performance indicators (KPIs) used to measure the reliability and efficiency of systems, processes, and equipment. Both metrics are important to risk management processes, providing measurable insights into potential risks and supporting risk mitigation strategies. MTTR and MTBF guide decisions regarding system design, maintenance practices, and redundancy or failover requirements.

Mean time between failures (MTBF) represents the expected lifetime of a product. The calculation for MTBF is the total operational time divided by the number of failures. For example, if you have 10 appliances that run for 50 hours and two of them fail, the MTBF is 250 hours/failure (10*50)/2.

Mean time to repair (MTTR) is a measure of the time taken to correct a fault so that the system is restored to full operation. This can also be described as mean time to replace or recover. MTTR is calculated as the total number of hours of unplanned maintenance divided by the number of failure incidents. This average value can be used to estimate whether a recovery time objective (RTO) is achievable.

A lower MTTR indicates quicker restoration of functionality, reducing downtime and potential disruptions to operations. This information helps allocate resources, prioritize maintenance activities, and optimize repair processes. MTBF identifies the average time between system or equipment failures. A higher MTBF suggests greater reliability and longer intervals between failures,

which can affect maintenance scheduling, spare part management, and overall system performance. Based on MTBF data, organizations can make decisions regarding maintenance strategies, equipment replacement, and investments in improving reliability.

Risk Identification and Assessment

Risk identification is fundamental to managing cybersecurity risks. It includes recognizing risks such as malware attacks, phishing attempts, insider threats, equipment failures, software vulnerabilities, and nontechnical risks like inadequate policies or training. Risk identification methods include vulnerability assessments, penetration testing, security audits, threat intelligence, and other methods. Risk identification is the foundation for risk assessment and management practices. Effective risk identification processes allow organizations to make informed decisions regarding resource allocation, risk mitigation strategies, and overall risk management practices.

Risk Management Strategies

Risk management strategies describe the proactive and systematic approaches used to identify, assess, prioritize, and mitigate risks to minimize their negative impacts.

Risk mitigation (or remediation) is the overall process of reducing exposure to or the effects of risk factors. A countermeasure that reduces exposure to a threat or vulnerability describes risk deterrence (or reduction) . Risk reduction refers to controls that can either make a risk incident less likely or less costly (or perhaps both). For example, if fire is a significant threat, a policy strictly controlling the use of flammable materials on-site reduces likelihood while a system of alarms and sprinklers reduces impact by (hopefully) containing any incident to a small area. Another example is off-site data backup, which provides a remediation option in the event of servers being destroyed by fire.

Avoidance means to stop the activity that is causing risk. For example, a company may develop an internally developed application for managing inventory and then try to sell it. During the sales process, the application may be discovered to have numerous security vulnerabilities that generate complaints and threats of legal action. The company may decide that the cost of maintaining the security of the software is not worth the revenue it generates and its development. Avoidance is infrequently a credible option.

Risk Transference

Transference (or sharing) means assigning risk to a third party, such as an insurance company. Specific cybersecurity insurance or cyber liability coverage protects against fines and liabilities arising from data breaches and attacks.

Note that in this sort of case it is relatively simple to transfer the obvious risks, but risks to the company's reputation remain. If a customer's credit card details are stolen because they used your unsecure e-commerce application, the customer won't care if you or a third party were nominally responsible for security. It is also unlikely that legal liabilities could be completely transferred in this way. For example, insurance terms are likely to require that best practice risk controls have been implemented.

It is not possible to eliminate risks, so a major objective of risk management is to determine an appropriate level of allowable risk. The concept of "allowable risk" varies greatly between organizations and is dependant on industry sector, leadership style, legal environment, and other factors.

Risk Acceptance

Risk acceptance (or tolerance) means that no countermeasures are put in place because the level of risk does not justify it.

A risk exception describes a situation where a risk cannot be mitigated using standard risk management practices or within a specified time frame due to financial, technical, or operational conditions. A risk exception formally recognizes the risk and seeks to identify alternate mitigating controls, if possible. Relevant stakeholders, such as risk managers or senior executives, must

approve all risk exceptions. Risk exceptions should be temporary and reviewed on an established time frame to determine whether the risk levels have changed or if the exception can be removed.

A risk exemption is a condition where risk can remain without mitigation, usually due to a strategic business decision. Risk exemptions are generally associated with situations where the cost of mitigating a risk outweighs its potential harm or can lead to significant strategic benefits when accepted. Similarly to risk exceptions, risk exemptions must be formally documented and approved by risk managers or senior executives and periodically reviewed using an established timetable.

The four risk responses are avoid, accept, mitigate, and transfer.

Where inherent risk is the risk before mitigation, residual risk is the likelihood and impact after specific mitigation, transference, or acceptance measures have been applied. Risk appetite is a strategic assessment of what level of residual risk is tolerable. Risk appetite is broad in scope. Where risk acceptance has the scope of a single system, risk appetite has a project- or institution-wide scope. Risk appetite is constrained by regulation and compliance.

12.1.3 Analyzing Risks (Lesson Video)

Transcript:

In this lesson, we're going to discuss risk analysis. Risk analysis is the practice of determining which identified threats are most pressing to an organization and then attaching a potential cost that can be expected if that threat occurs.

There are two general methods for analyzing and calculating risk. Quantitative analysis assigns real numbers to the costs of damages and countermeasures. It also assigns concrete probability percentages to risk occurrence. Qualitative analysis uses scenarios to identify risks and responses. Qualitative analysis is the more subjective of the two.
When calculating quantitatively, you want to consider a risk's impact and likelihood. What kind of damage can be done, and what's the chance of it actually happening? To determine the likelihood of a loss, you might want to use an annualized rate of occurrence, or ARO. An ARO is a collection of data gathered from historical records such as crime statistics, natural disasters, insurance payouts, and cyber incidence records.
To determine a loss's impact, you want to consider all associated costs if an unfortunate scenario were to occur. What would it cost to pay employees during the downtime? What would it cost to replace the impacted hardware or software? Would there be refunds or other costs associated with user retention?
Single loss expectancy, or SLE, is the amount of loss expected for any single successful asset attack. This is a monetary value that describes how much the incident will cost in terms of lost asset value.
Annual loss expectancy, or ALE, estimates the annual loss resulting from an incident. You calculate ALE by multiplying your SLE by your ARO. This tells you how much potential threats cost each year.
For example, if your asset loses $1,000 for each incident, and you expect an incident every four years, the annual cost for that asset would be $250.00.
Having a visual representation of risks can help stakeholders make better assessments. A risk register provides details of each known risk, including a risk's category, description, unique identification number, projected impact, likelihood of occurrence, and risk response plan. For better visualization, you could use this information to create a scatter plot that represents each risk's possible impact in relation to its overall probability.
A risk matrix, which is also known as a heat map, illustrates a loss's likelihood and impact. Notice that the colors in this example range from cool to hot. If an occurrence's possibility is rare and its impact moderate, it gets a low rating. If its likelihood is high, but the impact is insignificant, it might end up with a medium rating. And if the impact is severe and the likelihood almost certain, it might receive an extreme rating.
That's it for this lesson. In this lesson, we discussed risk analysis. We looked at calculating risks quantitatively and qualitatively, and we discussed a few ways to illustrate our findings.

12.1.4 Analyzing Risks Facts

Effective risk management practices involve systematically identifying, assessing, mitigating, and monitoring organizational risks. Audits provide an independent and objective evaluation of processes, controls, and compliance, ensuring adherence to standards and identifying gaps that pose risks. On the other hand, assessments help evaluate the effectiveness of risk management strategies, identify potential vulnerabilities, and prioritize mitigation efforts. By combining audits and assessments, organizations can comprehensively understand risks, implement appropriate controls, and continuously monitor and adapt their risk management strategies to protect against potential threats. These practices are essential for maintaining proactive and resilient security operations while ensuring compliance with legal mandates.

This lesson covers the following topics:

- Risk management processes
- Risk assessment
- Risk registers
- Risk threshold
- Key risk indicators
- Levels of risk appetite

Risk reporting

Risk Management Processes

For each business process and each threat, you must assess the degree of risk that exists. Calculating risk is complex, but the two main variables are likelihood and impact:

Likelihood is often used in qualitative analysis to describe the chance of a risk event happening subjectively. Likelihood is typically expressed using "low," "medium," and "high" or scored on a scale from 1 to 5.

Probability is a quantitative measure typically expressed as a numerical value between 0 and 1 or a percentage. Probability aims to precisely measure the chance of a risk event occurring based on statistical methods.

Impact is the severity of the risk if realized as a security incident. This may be determined by factors such as the value of the asset or the cost of disruption if the asset is compromised.

Risk management is complex and treated very differently in companies and institutions of different sizes and with different regulatory and compliance requirements. Most companies will institute enterprise risk management (ERM) policies and procedures based on frameworks such as NIST's Risk Management Framework (RMF) or ISO 31K. These legislative and framework compliance requirements are often formalized as a Risk and Control Self-Assessment (RCSA). An organization may also contract an external party to lead the process, referred to as a Risk and Control Assessment (RCA).

An RCSA is an internal process undertaken by stakeholders to identify risks and the effectiveness with which controls mitigate those risks. RCSAs are often performed through questionnaires and workshops with department managers. The outcome of an RCSA is a report. Up-to-date RCSA reports are critical to the external audit process.

Risk Assessment

Risk assessment is a core component of a cybersecurity program that evaluates previously identified risks to determine their potential impact on the organization. Risk assessment methodologies include ad hoc, recurring, one-time, or continuous:

Ad hoc risk assessments are conducted as needed, often in response to specific incidents, such as news of a new, actively exploited zero-day vulnerability or environmental changes such as system upgrades.

Recurring risk assessments are scheduled at regular intervals, such as annually, quarterly, or monthly, and can include audits, compliance checks, vulnerability scans, and other types of assessments.

One-time assessments are comprehensive evaluations carried out at a particular point in time, often during the implementation of a new system (or process) or to obtain an independent assessment of an organization's operational maturity.

Continuous risk assessments constantly evaluate risks and are supported by specialized tools that produce real-time data, such as agent-based vulnerability scanning platforms and intrusion detection systems. Different risk assessment methods are commonly combined to ensure effective identification and management of risk.

Quantitative risk assessment aims to assign concrete values to each risk factor. (Image © 123RF.com.)

Risk analysis describes the process of identifying and evaluating potential risks and the characteristics that define them. Risk analysis aims to understand the nature and scope of risks by examining their causes, consequences, and concerns.

Risk assessment is a systematic approach designed to estimate potential risk levels and their significance by interpreting data collected during risk analysis. Risk assessment considers the likelihood of an event occurring and the severity of its consequences. It may also involve prioritizing risks based on their potential impact and defining risk management strategies.

Quantitative Analysis

Quantitative risk analysis aims to assign concrete values to each risk factor.

Single Loss Expectancy (SLE) — The amount that would be lost in a single occurrence of the risk factor. This is determined by multiplying the value of the asset by an exposure factor (EF). EF is the percentage of the asset value that would be lost. For example, it may be determined that a tornado weather event will damage 40% of a building. The exposure factor in this case is 40% because only part of the asset is lost. If the building is worth $200,000, this event SLE is 200,000*0.4 or $80,000.

Annualized Loss Expectancy (ALE) — The amount that would be lost over the course of a year. This is determined by multiplying the SLE by the annualized rate of occurrence (ARO) . ARO describes the number of times in a year that an event occurs. In our previous (highly simplified) example, if it is anticipated that a tornado weather event will cause an impact twice per year, then the ARO is considered to be simply "2." The ARO is the cost of the event (SLE) multiplied by the number of times in a year it occurs. In the tornado example, SLE is $80,000, and ARO is 2, so the ALE is $160,000. This number is useful when considering different ways to protect the building from tornados. If it is known that tornados will have a $160,000 per year average cost, then this number can be used as a comparison when considering the cost of various protections.

It is important to realize that the value of an asset does not refer solely to its material value. The two principal additional considerations are direct costs associated with the asset being compromised (downtime) and consequent costs to intangible assets, such as the company's reputation. For example, a server may have a material cost of a few hundred dollars. If the server

were stolen, the costs incurred from being unable to do business until it can be recovered or replaced could run to thousands of dollars. In addition, the period of interruption during which orders cannot be taken or go unfulfilled may lead customers to seek alternative suppliers, potentially resulting in the loss of thousands of sales and goodwill.

The value of quantitative analysis is its ability to develop tangible numbers that reflect real money. Quantitative analysis helps to justify the costs of various controls. When analysts can associate cost savings with a control, it is easy to justify its expense. For example, it is easy to justify the money spent on a load balancer to eliminate losses from website downtime that exceeded the cost of the load balancer. Unfortunately, such direct and clear associations are uncommon!

The problem with quantitative risk assessment is that the process of determining and assigning these values is complex and time-consuming. The accuracy of the values assigned is also difficult to determine without historical data (often, it has to be based on subjective guesswork). However, over time and with experience, this approach can yield a detailed and sophisticated description of assets and risks and provide a sound basis for justifying and prioritizing security expenditure.

Qualitative Analysis

Qualitative risk analysis is a method used in risk management to assess risks based on subjective judgment and qualitative factors rather than precise numerical data. Qualitative risk analysis aims to provide a qualitative understanding of risks, their potential impact, and the likelihood of their occurrence. Often referred to as risk analysis using words, not numbers, this approach helps identify and prioritize intangible risks.

One of the benefits of qualitative risk analysis is its simplicity and ease of use. It does not require complex mathematical calculations or extensive data collection, making it a more accessible approach. It allows for a quick initial assessment of risks, enabling organizations to identify and focus on the most significant issues. Qualitative risk analysis frames risks by considering their causes, consequences, and potential interdependencies to improve risk communication and decision-making.

Qualitative risk analysis has some limitations. It is subjective in nature and heavily relies on expert judgment, which often introduces biases and inconsistencies if expert opinions differ. The lack of numerical data in qualitative risk analysis may make communicating risks to stakeholders who prefer quantitative information challenging. Despite these limitations, qualitative risk analysis is important because it provides a simplified description of risks and can help quickly draw attention to significant issues.

Inherent Risk

The result of a quantitative or qualitative analysis is a measure of inherent risk. Inherent risk is the level of risk before any type of mitigation has been attempted.

In theory, security controls or countermeasures could be introduced to address every risk factor. The difficulty is that security controls can be expensive, so it is important to balance the cost of the control with the cost associated with the risk. It is not possible to eliminate risk; rather, the aim is to mitigate risk factors to the point where the organization is exposed only to a level of risk it can tolerate. The overall status of risk management is referred to as risk posture. Risk posture shows which risk response options can be identified and prioritized. For example, an organization might identify the following priorities:

- Regulatory requirements to deploy security controls and make demonstrable efforts to reduce risk. Examples of legislation and regulations that mandate risk controls include SOX, HIPAA, Gramm-Leach-Bliley, the Homeland Security Act, PCI DSS regulations, and various personal data protection measures.
- High-value asset, regardless of the likelihood of the threat(s).
- Threats with high likelihood (that is, high ARO).
- Procedures, equipment, or software that increase the likelihood of threats (for example, legacy applications, lack of user training, old software versions, unpatched software, running unnecessary services, not having auditing procedures in place, and so on).

Heat Map

Another simple approach is the heat map or "traffic light" impact matrix. For each risk, a simple red, yellow, or green indicator can be put into each column to represent the severity of the risk, its likelihood, the cost of controls, and so on. This approach is simplistic but does give an immediate impression of where efforts should be concentrated to improve security.

Risk Factor	Impact	ARO	Cost of Controls	Overall Risk
Legacy Windows Clients				
Untrained Staff				
No Antivirus Software				

Traffic light impact grid.

FIPS 199 (nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.199.pdf) discusses how to apply security categorizations (SC) to information systems based on the impact that a breach of confidentiality, integrity, or availability would have on the organization as a whole. Potential impacts can be classified as the following:

Low — minor damage or loss to an asset or loss of performance (though essential functions remain operational).

Moderate — significant damage or loss to assets or performance.

High — major damage or loss or the inability to perform one or more essential functions.

Risk Registers

A risk register is a document showing the results of risk assessments in a comprehensible format. It includes information regarding risks, their severity, the associated owner of the risk, and all identified mitigation strategies. The register may include a heat map risk matrix (shown earlier) with columns for impact and likelihood ratings, date of identification, description, countermeasures, owner/route for escalation, and status.

Risk registers are also commonly depicted as scatterplot graphs, where impact and likelihood are each an axis, and the plot point is associated with a legend that includes more information about the nature of the plotted risk. A risk register should be shared among stakeholders (executives, department managers, and senior technicians) so that they understand the risks associated with the workflows that they manage.

Risk Threshold

Risk threshold defines the limits or levels of acceptable risk an organization is willing to tolerate. The risk threshold represents the boundaries within which risks are considered to be acceptable and manageable. Risk thresholds are based on various factors such as regulatory requirements, organizational objectives, stakeholder expectations, and the organization's risk appetite to help establish clear guidelines for decision-making. Organizations often define different risk thresholds for different types of risks based on their potential impact and criticality.

Key Risk Indicators

Key Risk Indicators (KRIs) are critical predictive indicators organizations use to monitor and predict potential risks. These metrics provide an early indication of increasing risk exposures in different areas of the organization. KRIs assess the potential impact and likelihood of various risks so leadership teams can take proactive steps to manage them effectively.

Using KRIs is closely associated with risk registers and risk management practices because KRIs provide the data needed to assess the likelihood and potential impact of each risk item tracked in a risk register. For example, a KRI may identify an increasing trend in system downtime due to IT operational issues that impact business operations. Risk managers handle this via a risk register and include details like potential impacts (lost productivity, customer dissatisfaction), mitigation steps (increasing IT resources, improving system redundancy), and the person or team responsible for managing these mitigations.

A risk owner refers to the individual responsible for managing a particular risk, including identifying and assessing the risk, implementing measures to mitigate it, monitoring the effectiveness of the measures, and taking corrective actions as warranted. The risk owner has a comprehensive understanding of the risk and its potential impacts and a thorough understanding of the measures needed to manage it. This role is often assigned to leadership team members with the authority to make decisions and the ability to allocate resources for risk mitigation. The risk owner also communicates information about the risk and its status to other stakeholders.

The risk appetite describes the level of risk that an organization is willing to accept. The organization's risk appetite is critical in determining which risks are added to a risk register and how they are prioritized. Risks are compared to the organization's risk appetite when identified and assessed. Risk tolerance describes the specific amount of variance an organization is willing to accept regarding measured risk levels and the established risk appetite. If a risk item's potential impact or likelihood exceeds the organization's risk tolerance, the risk is added to the risk register for appropriate management and monitoring. Risks that exceed the organization's risk tolerance by a large margin are generally prioritized and treated more urgently than other risks. In contrast, if a risk is near or slightly above the tolerance threshold, leadership teams may decide to accept it and monitor it closely.

Levels of Risk Appetite

Level	Description
Expansionary	An organization with an expansionary risk appetite is willing to take on higher levels of risk in the pursuit of high returns or aggressive growth. These organizations typically operate in rapidly evolving markets or industries and must take risks to remain competitive. Expansionary risk appetites are associated with organizations launching new products, entering new markets, or making major corporate acquisitions.
Conservative	An organization with a conservative risk appetite prioritizes risk avoidance. This type of organization takes a cautious approach to risks and prioritizes preserving cash, maintaining a good reputation, or ensuring regulatory compliance over pursuing aggressive growth.
Neutral	An organization with a neutral risk appetite balances expansionary and conservative approaches. It is willing to take on risks if they align with strategic objectives and can be managed effectively.

Risk Reporting

Risk reporting describes the methods used to communicate an organization's risk profile and the effectiveness of its risk management program. Effective risk reporting supports decision-making, highlights concerns, and ensures stakeholders understand the organization's risks. The content of risk reports must be relevant to its intended audience. For example, reports designed for board members must focus on strategic risks and the organization's overall risk appetite. Operational risk reports must include specific details regarding the factors contributing to risk and are appropriate for managers or technical employees.

Risk reports must also clearly convey recommended risk responses, such as accepting, mitigating, transferring, or avoiding the risk.

12.1.5 Business Continuity Planning (Lesson Video)

Transcript:

In this video, we're going to discuss business continuity planning. Natural disasters can cause catastrophic failures or extended disruptions--anything from a lightning strike that fries a few hardware components to an earthquake that destroys the entire building. Business continuity planning is the creation and validation of a plan for recovering and restoring an organization's mission-critical functions after a catastrophic disaster or extended disruption.

One important component of a business continuity plan is a business impact analysis, or BIA. A BIA focuses on the impact that losses will have on the organization. In a BIA, you identify critical functions--the operational functions that are necessary to conduct your business. Then you prioritize those critical business functions, calculate a timeframe for recovering them after a disaster, and estimate the tangible and intangible impact on the organization.

When you're creating this plan, you'll want to consider several parameters. The first is the recovery time objective, or RTO. The RTO states how much down time your organization can tolerate. It's another way of asking how long it'll take to get your system back up and running in the event of a disaster. Some data protection solutions have a very short recovery time and can bring your network back into operation relatively quickly. Others have a long RTO that lasts several days.

The second parameter is the recover point objective, or RPO. The RPO tells you two things. The first is the maximum length of time you can tolerate between backups. How old can restored data be? For example, if you run backups nightly, the maximum age of your data would be 24 hours because there would be a 24-hour window when no backups were taking place. If your organization only backs up data once a week, the maximum age of your data could be seven days. The longer the time between your backups, the older the data will be.

In an extreme example, if you back up once a week on Sunday and have a catastrophic failure on Friday, all the work done for five or six days would be lost. Most organizations can't tolerate that. But if you run your backups once a day, the maximum age becomes 24 hours. That sounds much better than five or six days, right?

The second thing RPO specifies is the number of backups from which to choose. Some disaster recovery solutions provide one backup. Other solutions may have multiple recovery points to select from. This depends on the data protection itself and how you've configured it.

Another thing to consider during your business impact analysis is the mean time to repair. MTTR identifies the average amount of time needed to repair a failed component or to restore operations. If you provide system access to other organizations, your service level agreement may specify how quickly you guarantee a fully functional and accessible system.

Similarly, the mean time between failures, or MTBF, identifies the average lifetime of a system or component--how long before the system wears out from age or use. Components should be replaced about the time the MTBF is reached.

In the event of a tornado, fire, or hurricane, your building will probably be unusable. Implementing site resiliency â€“ having an alternate location where you can pick up business where you left off â€“ could be extremely beneficial to your organization. There are three categories of site resiliency: hot sites, warm sites, and cold sites.

A hot site, or active backup model, is an alternate location that facilitates a fast and full recovery. This is, by far, the most effective backup option. It's also the most costly. In order to be fully operational within minutes, you need to have an alternate location. All servers, devices, cabling, and infrastructure need to be installed and ready to go. Hot sites often have a live connection to your main site with data copied across to the duplicate servers. This is called mirroring. If the active site fails, the hot site is fully operation immediately, with all the current data on the servers. This transition can occur within minutes or hours, so the downtime is minimal.

A warm site is partially configured. You have an alternate location. Some servers, some devices, some cabling, and some infrastructure is installed, but not everything. Data backups are available, but they aren't mirrored. Some recovery from backup tapes, the cloud, or other backup locations may be needed. This transition could occur within hours or days, so the downtime would be noticeable.

A warm site could be shared with another organization to lower its cost. This can be a good idea. Just make sure that the companies are in different areas so a natural disaster can't take them out at the same time.

Finally, there's a cold site. You do have an alternate location, but all equipment would need to be installed, and all data would need to be recovered. Downtime could extend for several weeks or even months.

No matter what type of site you decide to use, be sure to thoroughly document what you have and what you will need at the alternate location in the event of a disaster.
That's it for this lesson. In this video, we discussed business continuity plans, business impact analysis, and site resiliency, including hot, warm, and cold site architecture.

12.1.6 Business Continuity Planning Facts

Business continuity takes a broad approach, considering not only the continuity of critical functions but also the overall resilience and recovery of the entire organization. Business continuity planning includes the assessment of risks, the development of strategies to mitigate those risks, and the creation of plans to maintain or restore business operations in the face of various threats. This may involve addressing supply chain management, employee safety and communication, legal and regulatory compliance, and reputation management. Business continuity aims to ensure the long-term viability of an organization following a disruption, encompassing both immediate response and ongoing recovery efforts.

This lesson covers the following topics:

- Business continuity plan
- Continuity of operations
- Capacity planning
- Site considerations
- Cloud as disaster recovery (DR)
- Diversity and defense in depth
- Testing Resiliency
- Documentation

Business Continuity Plan

A business continuity plan (BCP) identifies appropriate disaster responses that maintain business operations during reduced or restricted infrastructure and resource capabilities. Also, a BCP identifies actions required to restore the business to normal operation. A business continuity plan is designed to ensure that critical business functions (CBF) can be performed when operations are disrupted.

A good plan documents all important decisions before the disaster strikes. When a disaster occurs, staff members simply need to follow the documented procedures. The process would be similar to the following:

- Identify the disaster, ensure the safety of personnel, and begin to implement recovery procedures.
- Implement short-term recovery mechanisms to bring mission-critical systems online.
- Stabilize operations by restoring supporting departments and functions.
- Implement measures to restore all functions to normal. Switch back from temporary measures to normal operating procedures. The order of restoration is defined in the BCP and then carried out in this last phase. A typical restoration order begins with the systems, databases, and applications that are most critical to the continued operation of the business. The order of restoration often varies significantly from one company to another.

Additional content in the BCP should include:

- Define processes for implementing, testing, and training team members. Team members should be representatives from all major parts of the corporation.
- After the plan has been created, conduct regular practices and training exercises to test portions of the plan. Revise the plan or training as necessary.
- As a BCP or DRP plan evolves, it is essential to collect and destroy all outdated copies of the plan as a new version of the plan is rolled out.
- Assign responsibility for ongoing maintenance of the BCP and DRP plans.

Succession planning is a process for identifying and developing internal people with the potential to fill future key positions within an organization. Succession planning:

- Increases the availability of experienced and capable employees that are prepared to assume specific roles as they become available.
- Ensures that the right competencies are recruited into the organization to be nurtured and developed over time to guarantee smooth transitions for future vacancies.
- Contrasts replacement planning, which focuses on identifying specific backup candidates for given positions.

Continuity of Operations

Continuity of operations (COOP) refers to the process of ensuring that an organization can maintain or quickly resume its critical functions in the event of a disruption, disaster, or crisis. COOP concepts and strategies aim to minimize downtime, protect essential resources, and maintain business resilience. Key elements of a COOP plan include identifying critical business functions, establishing priorities, and determining the resources needed to support these functions. Strategies often involve creating redundancy for IT systems and data, such as implementing off-site backups, failover systems, and disaster recovery solutions. Additionally, organizations may consider alternative work arrangements, such as remote work or co-location arrangements, to maintain operations during a crisis. Developing clear communication and decision-making protocols ensures that employees understand their roles and responsibilities during an emergency.

Regular testing and updating of continuity of operations plans (COOP) are crucial to ensure the organization can maintain essential functions during and after disruptive events. Realistic scenarios designed to simulate various disruptions, such as natural disasters, cyberattacks, or pandemics, must be used to assess the plan's effectiveness. Testing methods often include tabletop exercises, isolated functional tests, or full-scale drills. Each approach provides different levels of assurance and must therefore use pre-established evaluation criteria for measuring performance. In essence, COOP strategies focus on proactively preparing for disruptions, ensuring that organizations can continue to deliver essential services and minimize the impact of unforeseen events on their operations.

Backups

Backups play a critical role in the continuity of operations plans (COOP) by safeguarding against data loss and restoring systems and data in the event of disruptions. Regular testing verifies the integrity and effectiveness of backups. Testing backups helps ensure the backup process functions correctly by simulating various scenarios and allows organizations to identify any issues or gaps in the backup and recovery process. Testing backups validates the recoverability of critical systems and data, reducing the risk of data loss and minimizing downtime associated with disruptive events. Additionally, testing backups allows organizations to assess their recovery plans, evaluate the speed and efficiency of their backup systems, and ensure compliance with regulatory requirements. Inadequate backup processes can lead to extended downtime, critical data loss, financial losses, reputation damage, and noncompliance.

Relationship to Busines Continuity

Continuity of operations (COOP) and business continuity (BC) are closely related concepts that both focus on maintaining the ongoing functioning of an organization during and after a disruption, disaster, or crisis. However, they differ slightly in terms of their scope and primary objectives.

Continuity of operations primarily addresses the continuity of critical functions and services within an organization during an emergency or disaster. It often involves the development and implementation of strategies to maintain or restore essential operations, such as redundant IT systems, off-site backups, and disaster recovery solutions. COOP usually encompasses a shorter time frame, focusing on the immediate response to a disruption and the steps taken to resume critical functions as quickly as possible.

Continuity of operations is a component of the broader business continuity concept. Both are focused on maintaining the ongoing functioning of an organization during disruptions, but COOP is primarily concerned with the immediate response and restoration of critical functions, while business continuity encompasses a more comprehensive approach to ensure the overall resilience and recovery of the entire organization.

Capacity Planning

Capacity planning is a critical process in which organizations assess their current and future resource requirements to ensure they can efficiently meet their business objectives. This process involves evaluating and forecasting the necessary resources in terms of people, technology, and infrastructure to support anticipated growth, changes in demand, or other factors that may impact operations. For people, capacity planning considers the number of employees, their skill sets, and the potential need for additional training or hiring to meet future demands. This may involve evaluating workforce productivity, analyzing staffing levels, and identifying potential skills gaps. In terms of technology, capacity planning encompasses the assessment of hardware, software, and network resources required to support business operations, taking into account factors such as performance, scalability, and reliability.

Organizations must ensure they have the right technology resources in place to handle increasing workloads and support new applications or services. When it comes to infrastructure, capacity planning involves evaluating physical facilities, such as datacenters and office spaces, to determine whether they can accommodate projected growth and maintain business continuity. This may include considerations for power, cooling, and connectivity, as well as planning for potential expansion or relocation. Organizations use various capacity planning methods, including trend analysis, simulation modeling, and benchmarking, to help forecast their needs. Trend analysis examines historical data to identify patterns and trends in resource usage, demand, and performance. Organizations can forecast future resource requirements by understanding past patterns. This type of analysis can help identify potential bottlenecks or other areas that require attention. Simulation modeling leverages computer-based models to simulate real-world scenarios. Organizations can assess the impact of changes in demand, different resource allocation strategies, or system configurations to make informed decisions and optimize resource allocation to meet anticipated needs. Benchmarking requires a comparison of an organization's performance metrics against industry standards or best practices. Benchmarking provides a comparatively simple way to identify areas for improvement and establish performance targets. Ultimately, effective capacity planning allows organizations to optimize resource allocation, reduce costs, and minimize the risk of downtime or performance issues, ensuring they can continue to meet their business goals and maintain a competitive edge.

People Risks Associated with Capacity Planning

People risks associated with capacity planning may include insufficient staffing or skills gaps, leading to inadequate resource allocation or ineffective utilization. Lack of cross-training or succession planning can create dependency on specific individuals, increasing vulnerability to disruptions. Additionally, resistance to change, lack of employee engagement, or ineffective communication can hinder successful security operations.

Cross-Training — Requires employees to develop skills and knowledge outside their primary roles to mitigate the risk of relying heavily on specific individuals or teams. By cross-training employees, organizations can ensure that multiple individuals can perform critical tasks, reducing the dependence on a single employee or team. Cross-training promotes flexibility, resilience, and continuity within the workforce.

Remote Work Plans — Outline strategies for employees to work effectively outside the traditional office environment. Remote work plans define communication channels, technology requirements, and expectations for remote work arrangements. Established remote work plans allow organizations to seamlessly transition to remote operations, ensuring business continuity and minimizing the impact of disruptions.

Alternative Reporting Structures — Describe backup or temporary reporting relationships to reduce the risk associated with single points of failure in management or decision-making. Organizations can maintain operations and decision-making even if key personnel are unavailable by identifying interim individuals or teams.

Effective communication is paramount in reducing risk during disruptive events. Clear and timely communication channels ensure that employees, stakeholders, and customers receive accurate information, instructions, and updates. Clear communication helps manage expectations, reduce confusion, and facilitate coordinated responses. Communication fosters trust, promotes employee engagement, and ensures everyone is aligned with the organization's response plans. Additionally, communication plays a vital role in disseminating information about alternative work arrangements, changes in reporting structures, and other critical updates during a disruptive event.

Technologies and software associated with remote work:

Virtual Private Network (VPN) — Provides secure access to an organization's internal network and resources.

Remote Desktop Software — Allows remote access to computers or virtual desktops in the office or the tools used by the help desk and service-center teams to support employees.

Cloud-Based Tools — Platforms like Microsoft 365, Google Workspace, Dropbox, Slack, and many other popular tools enable remote team collaboration, document sharing, and communication.

Video Conferencing Software — Applications like Zoom, Microsoft Teams, or Webex facilitate virtual meetings, conference calls, and screen sharing.

Instant Messaging and Chat Tools — Applications like Slack, Microsoft Teams, or Discord enable real-time communication and quick collaboration.

Virtual Phone Systems — Cloud-based phone systems allow employees to make and receive calls remotely using their computers or mobile devices.

Project Management Tools — Platforms like Trello, Asana, or Jira assist in task management, project tracking, and team coordination.

Changes in Workforce Capacity

Layoffs can introduce a range of cybersecurity and physical risks to an organization, making it crucial to consider these factors within capacity planning efforts. Disgruntled employees may pose a significant cybersecurity risk, potentially engaging in unauthorized access or misuse of sensitive data and systems. Additionally, the loss of experienced employees may result in insufficient knowledge transfer to the remaining staff, leading to security gaps and misconfigurations. Furthermore, improper revocation of access to systems and data for laid-off employees can leave organizations vulnerable.

In terms of physical risks, departing employees may resort to theft or sabotage of physical assets or exploit their knowledge of safety protocols and procedures to compromise the organization's security. Unauthorized access to premises is another concern if access credentials are not revoked promptly.

Capacity planning is essential in managing these risks. It enables organizations to assess resource requirements and make strategic decisions about staffing levels and resource allocation. By incorporating potential layoffs into capacity planning, organizations can proactively prepare for the associated risks and minimize their impact. Implementing robust offboarding

procedures, ensuring proper knowledge transfer, and maintaining a strong security culture among remaining employees are crucial in mitigating risks. In conclusion, considering the potential risks associated with layoffs during capacity planning helps organizations maintain a secure environment and protect their valuable assets.

Other Risks Associated with Poor Capacity Planning

Poor capacity planning regarding technology and infrastructure can have significant consequences for an organization's cybersecurity and physical security. Overloaded systems resulting from inadequate capacity planning can increase susceptibility to crashes, failures, and denial of service (DoS) attacks. Additionally, limited resources may lead to performance degradation, potentially causing organizations to neglect essential security measures and updates. Failing to invest in the right security technologies or maintain the necessary infrastructure to protect against emerging threats leaves organizations more vulnerable to cyberattacks.

Physically, poor capacity planning may result in insufficient investment in security measures such as access control systems, surveillance cameras, or secure facilities, exposing organizations to unauthorized access or theft. Overlooking capacity requirements for power and cooling can cause overheating or power failures in datacenters, leading to hardware failures, data loss, or downtime. Furthermore, inadequate planning for future growth can limit an organization's ability to scale its operations, potentially affecting its responsiveness to security incidents or implementation of new security measures.

In contrast, overestimating capacity needs during capacity planning can negatively impact an organization. Increased costs from unnecessary expenses on technology, infrastructure, and personnel, strain budgets and divert funds from other critical areas. Inefficient resource utilization can lead to low utilization rates, which can negatively affect the return on investment (ROI) and overall operational effectiveness. Overestimating capacity needs can contribute to higher energy consumption, driving up costs and increasing the organization's carbon footprint and environmental impact.

Deploying more resources than necessary can introduce increased complexity in managing and maintaining technology and infrastructure. This added complexity could create challenges for IT teams, making it more difficult to identify and resolve issues or optimize performance. Additionally, the opportunity cost of investing in excess capacity can be significant, as resources may be diverted from other essential projects or initiatives, potentially hindering innovation or market growth.

To avoid these potential problems, organizations must strive for a balanced approach to capacity planning, considering current and future needs while remaining flexible and adaptable to changing circumstances. Regularly reviewing and updating capacity plans, along with employing techniques such as monitoring, forecasting, and resource scaling, can help organizations optimize resource allocation and mitigate the risks associated with overestimating capacity needs.

Site Considerations

Enterprise environments often provision resiliency at the site level. An alternate processing or recovery site is a location that can provide the same (or similar) level of service. An alternate processing site might always be available and in use, while a recovery site might take longer to set up or only be used in an emergency.

Operations are designed to failover to the new site until the previous site can be brought back online. Failover is a technique that ensures a redundant component, device, application, or site can quickly and efficiently take over the functionality of an asset that has failed. For example, load balancers provide failover in the event that one or more servers or sites behind the load balancer are down or are taking too long to respond. Once the load balancer detects this, it will redirect inbound traffic to an alternate processing server or site. Thus, redundant servers in the load balancer pool ensure there is no or minimal interruption of service.

Site resiliency is described as hot, warm, or cold:

- A hot site can failover almost immediately. It generally means the site is within the organization's ownership and ready to deploy. For example, a hot site could consist of a building with operational computer equipment kept updated with a live data set. Hot sites often have a live connection to your main site with data copied across to the duplicate servers. This is the most costly option.

A warm site could be similar, but with the requirement that the latest data set needs to be loaded.

A cold site takes longer to set up. A cold site may be an empty building with a lease agreement in place to install whatever equipment is required when necessary.

Providing redundancy on this scale is generally very complicated and expensive. Another issue is that creating duplicate sites and data also doubles (or more) the complexity of securing the resources. The same security procedures must apply to redundant sites, spare systems, and backup data as apply to the main copy.

Geographic dispersion refers to the distribution of recovery sites across different geographic locations for disaster recovery (DR) purposes. The concept aims to ensure that recovery sites are located far enough apart to minimize the impact of regional disasters, such as natural calamities or localized incidents. By strategically dispersing recovery sites, organizations can ensure the resilience of their recovery strategies and reduce the risk of a single catastrophic event affecting all business operations.

Cloud as Disaster Recovery (DR)

Several factors drive organizations to use cloud services for datacenter or site redundancy. Cost efficiency plays a significant role, as cloud providers offer more affordable redundancy and backup options due to their economies of scale.

"Economies of scale" is a concept that refers to the cost advantages that businesses can achieve when they increase production and output. Essentially, the more a company produces, the cheaper it becomes to deliver those products.

The scalability of cloud services allows businesses to incorporate redundant capabilities without over-provisioning resources. For example, geographic diversity provided by cloud providers helps protect against regional outages or disasters, but incorporating this type of redundancy for a private organization would be cost-prohibitive and unwarranted.

Faster deployment of capabilities when using cloud platforms enables organizations to set up and deploy redundant systems quickly, far more rapidly than could be done when building infrastructure from scratch.

Simplified management is another critical factor, with cloud providers offering tools and services that reduce the complexity of managing redundant infrastructure. Improved security and compliance are also important considerations, as cloud providers invest heavily in these areas, helping organizations meet data protection and redundancy regulatory requirements.

Diversity and Defense in Depth

Platform diversity is a concept in cybersecurity that refers to using multiple technologies, operating systems, and hardware or software components within an organization's infrastructure. By incorporating a variety of platforms, businesses can reduce the risk of a single vulnerability or attack affecting their entire infrastructure. This approach is important for cybersecurity operations, as it helps create a more resilient environment that can better withstand cyber threats.

When an organization relies on a single technology or platform, an attacker who discovers a vulnerability can potentially compromise the entire system. However, with platform diversity, even if one component is compromised, other parts of the system remain secure, limiting the potential damage. Furthermore, a diverse technology landscape can make it more challenging for threat actors to navigate, as they must be familiar with multiple platforms and exploit techniques. In this way, platform diversity deters potential attackers and contributes to the overall robustness of an organization's cybersecurity posture.

Testing Resiliency

Testing system resilience and incident response effectiveness are crucial for organizations to recover from disruptions and maintain business continuity. By conducting various tests, organizations can identify potential vulnerabilities, evaluate the efficiency of their recovery strategies, and improve their overall preparedness for real-life incidents.

Tabletop Exercises involve teams discussing and working through hypothetical scenarios to assess their response plans and decision-making processes. These exercises help identify knowledge, communication, and coordination gaps, ultimately strengthening the organization's incident response capabilities. For example, a tabletop exercise might simulate a ransomware attack to test how well the organization's IT and management teams collaborate to mitigate the threat and restore operations.

Failover Tests involve intentionally causing the failure of a primary system or component to evaluate the automatic transfer of operations to a secondary, redundant system. These tests ensure backup systems can seamlessly take over during an actual incident, minimizing downtime and data loss. For example, a failover test could involve simulating the failure of a primary database server to verify that a standby server can successfully assume its role and maintain service continuity.

Simulations are controlled experiments replicating real-world scenarios, allowing organizations to assess their incident response processes and system resilience under realistic conditions. These tests can reveal potential bottlenecks, inefficiencies, or vulnerabilities that might not be apparent in less complex tests. For instance, a simulation might involve a cyberattack targeting the organization's network infrastructure to evaluate the effectiveness of security measures and the ability to detect, contain, and remediate the threat.

Parallel Processing Tests involve running primary and backup systems simultaneously to validate the functionality and performance of backup systems without disrupting normal operations. These tests help organizations ensure their backup systems can handle the same workload as primary systems during an incident. For example, an organization might run parallel processing tests to verify that a backup datacenter can manage the same traffic and processing demand as the primary datacenter in an outage.

Failure to perform tests such as tabletop exercises, failover tests, simulations, and parallel processing can expose organizations to significant risks. With these tests, organizations can recognize potential vulnerabilities and weaknesses in their incident response plans and system resilience designs and use the results of tests to improve existing plans. In a real-life disruption or cyberattack, untested systems and response procedures may fail to perform as expected, leading to extended downtime, data loss, and reputational damage. Moreover, unprepared organizations may face increased costs related to incident recovery and mitigation and potential regulatory penalties for failing to meet industry standards and compliance requirements. Ultimately, failure to implement these tests can leave organizations inadequately prepared for crises, undermining their ability to maintain business continuity and protect valuable assets.

Documentation

Business continuity documentation practices cover planning, implementation, and evaluation. Documentation supports the testing process. Documentation includes test plans outlining the objectives, scope, and methods of tests and the roles and responsibilities of individuals involved. Test scripts (or scenarios) provide step-by-step instructions for performing the tests, and test results identify strengths and weaknesses of the business continuity plan and the technical capabilities supporting it. Documentation is the foundation for clear communication and reporting of activities. It provides a common reference point for those involved in business continuity testing and facilitates effective communication with management, executive teams, and other relevant stakeholders. Third-party assessments and certifications offer an objective and independent evaluation of an organization's testing practices. Third-party assessments and certifications offer objective evaluation, compliance verification, validation of testing effectiveness, industry recognition, and recommendations for continuous improvement. Examples of third-party evaluations include assessments performed in alignment with ISO 22301, PCI DSS, and SOC 2.

12.1.7 Practice Questions (Section Quiz)

q_risk_types_tol_avoidance_secp8

A company identifies a potential security risk associated with the implementation of a new system. However, after assessing the risk, the company decides not to implement any measures to address this specific risk.

Which of the following risk management strategies is the company employing?

Answers:

***Avoidance**

Transference

Mitigation

Exemption

Explanation:

The avoidance strategy means that the organization will not proceed with actions likely to trigger the risk.

Transferring risk means shifting the risk to a third party, typically through insurance or outsourcing. This is not what the company is doing in this scenario.

Mitigation involves taking steps to reduce the severity or likelihood of the risk. In this scenario, the company is not implementing any measures to address the risk, so they are not mitigating it.

Exemption is generally a formal process documented and approved by risk managers or senior executives. The company chooses not to address the risk without indication of any formal exemption process, which is more aligned with risk avoidance.

q_risk_types_tol_inherent_risk_secp8

A small engineering company wants to run a business analysis. It hired a consulting firm to better understand the underlying components, including the result of quantitative or qualitative risk analysis.

Which of the following values is MOST beneficial to the company in this situation?

Answers:

***Inherent risk**

Annualized loss expentancy (ALE)

Risk factors

Single loss expentency (SLE)

Explanation:

The result of quantitative or qualitative analysis is a measure of inherent risk. Inherent risk is the level of risk before implementing mitigation.

ALE is the amount a company would lose over the course of a year. The calculation for ALE is the multiplication of SLE by the annualized rate of occurrence (ARO).

The SLE is the amount a company would lose in a single risk factor occurrence. The calculation for SLE is the multiplication of the asset value by an exposure factor (EF).

A risk factor is an item used as a risk input during quantitative or qualitative analysis.

q_risk_types_tol_mttr_01_secp8

A company experiences a significant system failure that leads to service interruption. The IT department works to restore the system and documents the duration it takes to fix the problem.

This recorded duration is primarily indicative of which of the following options?

Answers:

***Mean time to repair (MTTR)**

Risk assessment

Mean time between failures (MTBF)

Risk tolerance

Explanation:

Mean time to repair (MTTR) measures the average time required to repair a failed component or system.

Risk assessment involves identifying, evaluating, and prioritizing risks. This process doesn't primarily involve recording the duration of system repairs.

Mean time between failures (MTBF) measures the anticipated time between system failures during operation. This metric doesn't primarily concern the time taken to repair a system after it has failed.

Risk tolerance refers to the amount of risk that an organization is willing to accept in its operations. It does not primarily involve the duration of system repairs.

q_risk_types_tol_mttr_02_secp8

A risk manager for a company providing IT support services conducts a business impact analysis (BIA) and identifies a mission essential function (MEF) that relies on a server with a mean time between failures (MTBF) of 2,500 hours and a mean time to repair (MTTR) of 4 hours.

Given a maximum tolerable downtime (MTD) of 24 hours and a recovery time objective (RTO) of 6 hours for this function, what should the risk manager prioritize in the risk management strategy?

Answers:

Improving the MTBF of the server.

***Reducing the MTTR of the server.**

Increasing the MTD for the function.

Extending the RTO for the function.

Explanation:

Reducing the MTTR is crucial. With an RTO of 6 hours and an MTTR of 4 hours, other activities have limited time, such as recovery during the maximum tolerable downtime. Lowering the MTTR provides more time for those activities.

The current MTBF of 2,500 hours is substantial and not an immediate concern considering the MTD of 24 hours.

Increasing the MTD might not be feasible or desirable. The risk manager sets the MTD based on the disruption the business can tolerate.

Extending the RTO is not the best option. The current RTO of 6 hours seems adequate, considering the MTTR of 4 hours. The focus should be on restoring the server within the RTO.

q_risk_types_tol_mttr_mtbf_01_secp8

A company is reviewing its system reliability metrics. It needs to know the average time the system operates without failure and the average time it takes to repair a system when it fails.

Which of the following pairs of metrics should the company focus on to meet its needs?

Answers:

***Mean time between failures (MTBF) and mean time to repair (MTTR)**

Recovery time objective (RTO) and recovery point objective (RPO)

Risk tolerance and risk appetite

Risk assessment and risk analysis

Explanation:

MTBF measures the average time a system or component expects to operate without failure during typical operation. MTTR is the average time required to repair a failed component or device.

RTO and RPO are measures of the acceptable amount of data loss or downtime in case of a disaster, not measurements of system reliability or repair times.

Risk tolerance and risk appetite are not terms used to describe the metrics related to system failure and recovery.

Risk assessment and risk analysis are processes involved in identifying and evaluating risks and do not provide measurements for system reliability or repair times.

q_risk_types_tol_mttr_mtbf_02_secp8

A manufacturing organization identifies its server maintenance and repair process as a mission-essential function. The company experienced three server failures in the last year, each failure taking approximately six hours to repair and restore operations.

Given the company's performance metrics and assuming all failures occurred during operational hours, what are the MTBF and MTTR for the organization's server maintenance and repair process?

Answers:

MTBF: 2,000 hours/failure, MTTR: 6 hours

MTBF: 1,000 hours/failure, MTTR: 2 hours

***MTBF: 2,920 hours/failure, MTTR: 6 hours**

MTBF: 1,460 hours/failure, MTTR: 18 hours

Explanation:

The MTTR (mean time to repair) is 6 hours, which is the time it took to repair and restore operations for each failure. The MTBF (mean time between failures) is 2,920 hours/failure, calculated by dividing the total operational time of 8,760 hours per year by 3 failures.

Incorrect MTTR answers include 2 hours and 18 hours, as it took 6 hours to restore operations for each failure.

Incorrect MTBF answers include 2,000, 1,000, and 1,460 hours/failure. The correct MTBF should be 2,920 hours/failure, calculated by dividing the total operational time of 8,760 hours per year by 3 failures.

q_risk_types_tol_mttr_mtbf_03_secp8

Which of the following statements accurately distinguish the differences between mean time to repair (MTTR) and mean time between failures (MTBF)?

Answers:

***MTTR describes how long it would take to bring equipment back into operation, while the MTBF refers to how long equipment will last until it is no longer operational.**

MTTR refers to how long equipment will last until it is no longer operational, while the MTBF describes how long it would take to bring equipment back into operation.

MTTR is the amount of data loss that a system can sustain, measured in time, while MTBF is the longest period of time that a business function outage may occur without causing irrecoverable business failure.

MTBF is the amount of data loss that a system can sustain, measured in time, while MTTR is the longest period of time that a business function outage may occur for without causing irrecoverable business failure.

Explanation:

Mean time to repair (MTTR) is a measure of the time taken to correct a fault so that the company can bring a system back to full operation. We can also describe this as the mean time to replace or recover.

Mean time between failures (MTBF) represents the expected lifetime of a product.

The recovery point objective (RPO) is the amount of data loss that a system can sustain, measured in time.

Maximum tolerable downtime (MTD) is the longest period of time that a business function outage may occur without causing irrecoverable business failure.

q_risk_types_tol_mttr_mtbf_04_secp8

Which of the following statements accurately distinguishes the similarities between mean time to repair (MTTR) and mean time between failures (MTBF)?

Answers:

***Both MTTR and MTBF can determine the amount of asset redundancy a system should have.**

Both MTTR and MTBF are for non-repairable assets.

The MTBF and MTTR calculations are the same.

MTTR is the median time it takes to correct a fault, while MTBF is the average time it takes to correct a fault.

Explanation:

The similarity is that MTTR and MTBF can determine the amount of asset redundancy a system should have. A redundant system can fail over to another asset if there is a fault and continue to operate normally.

MTTR is for repairable assets, while MTBF is meant for non-repairable assets.

The calculation for MTBF is the total time divided by the number of failures. The calculation for MTTF is the total time divided by the number of devices.

MTTR is the average measure of the time taken to correct a fault so that the system restores to full operation. MTBF is the average time between failures.

q_risk_types_tol_percentage_secp8

An organization is expanding its operations into a new region with unfamiliar regulatory requirements. The risk management team conducts a thorough risk assessment and identifies a need for robust controls to ensure compliance.

Which of the following would be the MOST effective metric for tracking regulatory compliance risk in this situation?

Answers:

***The percentage of employees who have undergone compliance training.**

The total revenue generated from the new region.

The number of identified market competitors.

The frequency of audits conducted by the regulatory authority.

Explanation:

The percentage of employees who have undergone compliance training is the most relevant for managing regulatory compliance risk. Training employees on the new regulatory requirements ensure that they understand the rules and can act in compliance with them.

While revenue is an important business metric, it does not directly measure regulatory compliance risk.

The number of competitors may provide insights into market competition, but it does not directly measure regulatory compliance risk.

While audit frequency is a measure of regulatory scrutiny, the primary control for ensuring compliance lies in the organization's internal practices, not in the frequency of external audits.

q_risk_types_tol_qualitative_quantitative_secp8

A cybersecurity consultant is assessing risks for a new e-commerce website. The consultant identifies potential risks, evaluates their impact and likelihood, and considers the organization's ability to mitigate them.

Which risk assessment methodology is the consultant MOST likely using?

Answers:

Qualitative

Quantitative

***Qualitative and quantitative**

Ad hoc

Explanation:

The consultant considers both qualitative and quantitative factors. Qualitative factors include the risk's impact and likelihood of occurring, while quantitative factors would be the ability to mitigate risks in the risk assessment.

The consultant considers qualitative factors, such as impact and likelihood. However, the consultant must also consider the organization's ability to mitigate risks involving quantitative factors, such as costs.

The consultant must also consider qualitative factors, like impact and likelihood, and not exclusively quantitative factors.

Consultants conduct ad hoc assessments as needed, often responding to specific incidents. However, the scenario does not suggest or mention that this assessment is in response to a specific incident.

q_risk_types_tol_quantitative_secp8

The IT department at a multinational organization is evaluating potential risks associated with implementing a new network infrastructure. This includes identifying potential vulnerabilities, estimating potential downtime, and assessing the financial impact of potential cyberattacks.

Which type of risk assessment BEST suits the organization's requirements?

Answers:

***Quantitative risk assessment**

Qualitative risk assessment

Ad hoc risk assessment

Continuous risk assessment

Explanation:

Quantitative risk assessment allows the team to measure risk in terms of loss and occurrence by using numerical or measurable data.

While the qualitative risk assessment method categorizes risks into levels, it does not provide the exact numerical data the IT team needs for their assessment.

Ad hoc risk assessment is not the first choice as this type of assessment is usually carried out in response to an immediate threat or incident and might not provide a comprehensive review of all the potential risks.

While continuous risk assessment emphasizes ongoing evaluation and updating of the risk assessment process, it's not inherently quantitative or qualitative. In this scenario, the team is specifically looking for measurable data.

q_risk_types_tol_regulatory_compliance_secp8

A multinational company is looking to ensure that its global operations meet all the relevant legal requirements to avoid fines, penalties, and potential loss of reputation.

What type of process is the company planning to implement?

Answers:

***Regulatory compliance**

Risk threshold

Risk management

Audit

Explanation:

A regulatory process refers to a set of procedures to ensure that a company's operations meet the legal and regulatory requirements in their jurisdictions.

The risk threshold is the point when risks become unacceptable. The company is looking for regulatory compliance, which is not a risk to ignore in a risk threshold.

While risk management is a broader process that includes identifying, assessing, and prioritizing risks, it is not specific to adhering to legal and regulatory requirements.

An audit is a process of reviewing an organization's operations, procedures, or financials. While it may check compliance with regulations, it is different from the process of adhering to them.

q_risk_types_tol_residual_secp8

Your company has developed and implemented countermeasures for the greatest risks to their assets. However, there is still some risk left.

What is the remaining risk called?

Answers:

Exposure

Risk

Loss

***Residual risk**

Explanation:

Residual risk is the portion of risk that remains after the implementation of a countermeasure. There is almost always some residual risk.

Exposure is the vulnerability of losses from a threat agent, and risk is the likelihood of a vulnerability being exploited.

A loss is the real damages to an asset that reduces its confidentiality, integrity, or availability.

Risk is a generic term referring to risk management, which involves identifying potential issues, assessing their potential impact on the organization, and implementing controls to mitigate them.

q_risk_types_tol_risk_acceptance_secp8

A new company implements a data center that will hold proprietary data that is output from a daily workflow. As the company has not received any funding, no risk controls are in place.

How does the company approach risk during operations?

Answers:

***Risk acceptance**

Risk avoidance

Risk mitigation

Risk transference

Explanation:

Risk acceptance means that the company did not put in place countermeasures, either because the level of risk does not justify the cost or because there will be an unavoidable delay before the company can deploy the countermeasures.

Risk avoidance means that a company can stop an activity that is risk-bearing. The discontinuation of a defective product is a form of risk avoidance.

Risk mitigation is the overall process of reducing exposure to the effects of risk factors. An example is deploying a countermeasure that reduces exposure to a threat or vulnerability.

Risk transference means assigning risk to a third party, such as an insurance company or a contract with a supplier that defines liabilities.

q_risk_types_tol_risk_appetite_secp8

A company is looking to expand its business into new markets despite associated risks. It prepares to accept higher risks for potentially higher returns.

Which of the following approaches BEST meets the company's risk management approach parameters?

Answers:

***Risk appetite**

Risk threshold

Risk tolerance

Risk mitigation

Explanation:

Risk appetite refers to the level of risk an organization is willing to accept in pursuit of its objectives. The company has an "appetite" for a certain amount of risk.

The risk threshold is the point when risks become unacceptable. They're looking at the company's willingness to accept risk, not the point at which the risk becomes unacceptable.

Risk intolerance is an organization's unwillingness to accept risk. The company is willing to accept higher risks for potentially higher returns, making it tolerant of some level of risk.

Risk mitigation is the process of developing actions and plans to decrease the risks to an organization. The question is asking about the company's approach to accepting risks.

q_risk_types_tol_risk_assessment_01_secp8

A global corporation assesses risk appetite and how risks in various regions could influence mission-critical operations. It is assessing compliance with local laws and licensing requirements to prevent financial risk or resolve security risks, changing the risk posture, and implementing risk controls to compensate.

What type of assessment is the team performing?

Answers:

***Risk control assessment**

Site risk assessment

Vulnerability assessment

Penetration testing

Explanation:

Risk and control self-assessment (RCSA) is the method by which companies evaluate and analyze the operational risks and the efficacy of the controls used to manage them.

A site risk assessment evaluates exposure to safety concerns related to area-related risks, such as disaster risk, utility disruption, or health and safety concerns.

A vulnerability assessment is a comprehensive analysis of the security vulnerabilities of an information system. Automated vulnerability software, such as Nessus, is sometimes used.

Penetration testing is a simulated cyberattack on a computer system to discover any exploitable vulnerabilities.

q_risk_types_tol_risk_assessment_02_secp8

The security team at a company is adopting a cybersecurity framework to standardize its security measures across different departments. The team lead wants to ensure that the selected framework encompasses all the critical aspects of cybersecurity.

What should the security team lead ensure the cybersecurity framework covers to provide a comprehensive security posture?

Answers:

Procedures for incident response

Threat intelligence and event correlation

The technical controls and access management

***Risk assessment, incident response, access control, awareness, and training**

Explanation:

Risk assessment, incident response, access control, awareness, and training are the vital components of a comprehensive cybersecurity framework. They ensure that well-rounded security measures are effective against a variety of threats.

Although incident response is a critical aspect of cybersecurity, a comprehensive cybersecurity framework must cover much more than just incident response.

While these components are essential for proactive threat management, a comprehensive cybersecurity framework requires more components, including risk assessment, incident response, access control, and training.

Even though these components are important for securing the system, a comprehensive framework should also include risk assessment, incident response, awareness, and training for a well-rounded security posture.

q_risk_types_tol_risk_impact_01_secp8

A company is evaluating the potential outcomes of a certain risk event. It estimates that if the event occurs, it could lead to a financial loss measured in dollars.

Which of the following outcomes can the company conclude in this scenario?

Answers:

***Impact**

Risk tolerance

Exposure factor

Annualized loss expectancy (ALE)

Explanation:

The impact is the potential loss incurred if a risk event occurs. The financial loss that could result from the risk event represents the impact of that event.

Risk tolerance refers to the amount of risk that an organization is willing to accept. The scenario did not define the amount of risk it will accept.

The exposure factor is the percentage of loss an organization would experience during specific asset violations or damage. The measurement described is dollars and not a percentage.

Annualized loss expectancy (ALE) is the potential yearly cost of all instances of a specific risk. The scenario does not provide information about how often the risk event might occur annually.

q_risk_types_tol_risk_impact_02_secp8

A company has recently identified a potential risk involving its inventory management software. The company's risk manager has conducted an initial assessment and listed the risk in the risk register.

As a result, the manager must decide on a suitable response to this risk.

Which of the following is the MOST appropriate step for the risk manager?

Answers:

Identify and assess the potential vulnerabilities and threats associated with the risk

Identify the mission essential functions of the company and assess the risk's impact on them

***Determine the likelihood and impact of the risk on the company's operations**

Update the risk register with the mitigation strategies and inform the stakeholders

Explanation:

Determine the likelihood and impact of the risk on the company's operations before deciding on the appropriate risk response. The risk manager must understand these factors to determine whether to accept, mitigate, transfer, or avoid the risk.

The risk manager identifies and assesses vulnerabilities and threats during the risk identification process before deciding on the risk response.

The risk manager should have identified mission essential functions during the initial risk identification and assessment process.

The risk manager updates the risk register with the mitigation strategies, informing the stakeholders as the final step in the risk management process.

q_risk_types_tol_risk_tolerance_secp8

A company is willing to accept certain risks associated with a new project as it believes the potential rewards outweigh the potential downsides. However, it will not accept risks that jeopardize its core business functions.

What does this decision about the company's risk management approach primarily represent?

Answers:

***Risk tolerance**

Risk threshold

Risk appetite

Risk transference

Explanation:

Risk tolerance refers to the degree of variability in investment returns an organization is willing to withstand. The company's willingness to accept certain risks for a new project, but not those that could endanger core business functions, is an example of defining its risk tolerance.

The risk threshold is the specific point at which risk becomes unacceptable and requires action.

Risk appetite is the amount of risk an organization is willing to seek or accept to achieve its long-term objectives. However, the question is more focused on the company's risk tolerance.

Risk transference involves shifting the risk to another party, typically through insurance or outsourcing. The scenario does not involve transferring risk to another party.

q_risk_types_tol_risk_transference_01_secp8

A technology company experiences several security vulnerabilities with its online application, leading to customer complaints and legal threats. In response, the board of directors is considering outsourcing the maintenance of the application to a third party.

Which risk management strategy is the company primarily implementing?

Answers:

Risk avoidance

Risk acceptance

Risk mitigation

***Risk transference**

Explanation:

The company is transferring the risk to a third party by outsourcing the maintenance of the application, which is an example of risk transference.

Risk avoidance involves halting the risky activity. In this scenario, the company is not stopping its operations but shifting the maintenance to a third party.

Risk acceptance means the company has not deployed countermeasures because it has deemed the risk level acceptable. This is different from the company actively seeking a solution, indicating it is not accepting the risk.

Mitigation involves implementing controls to lessen the likelihood or impact of risk. Here, the company is not implementing controls but transferring the task to a third party.

q_risk_types_tol_risk_transference_02_secp8

A company relies on a legacy system for a core business process, but the system's vendor no longer provides security patches. To manage the risks associated with this system, the company decides to contract a third-party provider to take on this risk.

What strategy is the company using?

Answers:

***Risk transference**

Risk acceptance

Risk mitigation

Patch management

Explanation:

The company is transferring the risk associated with the inability to patch the legacy system to a third-party provider.

Risk acceptance involves accepting the risk without taking further steps to reduce its impact. Since the company is taking steps to manage the risk, this option does not apply.

Risk mitigation refers to efforts to reduce the impact of a risk. Contracting a third-party provider to manage the risk is not a risk mitigation strategy but a risk transference strategy.

Patch management is the process of managing updates for software applications. However, in this scenario, the company cannot patch the system, so patch management is not the strategy used.

q_risk_types_tol_risk_transference_03_secp8

A business utilizes an outdated system essential to its operations, but the supplier no longer offers updates for security flaws.

In response to the vulnerabilities linked to this system, the business chooses to engage an external provider to assume responsibility for this risk.

What approach is the business employing to handle this situation?

Answers:

***Risk transference**

Risk acceptance

Risk mitigation

Patch management

Explanation:

The company is transferring the risk associated with the inability to patch the legacy system to a third-party provider.

Risk acceptance involves accepting the risk without taking further steps to reduce its impact. Since the company is taking steps to manage the risk, this option does not apply.

Risk mitigation refers to efforts to reduce the impact of a risk. Contracting a third-party provider to manage the risk is not a risk mitigation strategy but a risk transference strategy.

Patch management is the process of managing updates for software applications. However, in this scenario, the company cannot patch the system, so patch management is not the strategy used.

q_risk_types_tol_rpo_01_secp8

An organization is continuously backing up its data to ensure minimum data loss in case of a system failure. It is trying to decide the maximum age of files that require recovery from backup storage for normal operations to resume after a failure.

What concept should the organization consider to meet their needs?

Answers:

***Recovery point objective (RPO)**

Single loss expectancy (SLE)

Recovery time objective (RTO)

Annualized loss expectancy (ALE)

Explanation:

A recovery point objective (RPO) is the maximum acceptable age of data that an organization can tolerate losing. When an organization continuously backs up its data, it is typically aiming to keep its RPO as low as possible.

Single loss expectancy (SLE) is a term used in risk assessment to represent the expected loss for any single event. It does not directly relate to data backup frequency or data age considerations.

While recovery time objective (RTO) and RPO are both important disaster recovery considerations, RTO does not determine the acceptable age of data for recovery.

Annualized loss expectancy (ALE) is not a measure related to the age of data that needs recovery after a failure.

q_risk_types_tol_rpo_02_secp8

A technology company implements a backup strategy to mitigate data loss in case of a system crash. The strategy focuses on defining the maximum duration of data to retrieve from the backup storage to ensure business continuity after a system crash.

Which principle should the company applying to meet their needs?

Answers:

***Recovery point objective (RPO)**

Single loss expectancy (SLE)

Recovery time objective (RTO)

Annualized loss expectancy (ALE)

Explanation:

A recovery point objective (RPO) is the maximum acceptable age of data that an organization can tolerate losing. When an organization continuously backs up its data, it is typically aiming to keep its RPO as low as possible.

Single loss expectancy (SLE) is a term used in risk assessment to represent the expected loss for any single event. It does not directly relate to data backup frequency or data age considerations.

While recovery time objective (RTO) and RPO are both important disaster recovery considerations, RTO does not determine the acceptable age of data for recovery.

Annualized loss expectancy (ALE) is not a measure related to the age of data that needs recovery after a failure.

q_risk_types_tol_rto_01_secp8

An organization performs a business impact analysis to identify potential effects of business interruptions. It is trying to identify the maximum acceptable time its key business process can be down before it severely impacts operations.

What is the organization attempting to determine?

Answers:

***Recovery time objective (RTO)**

Mean time to repair (MTTR)

Mean time between failures (MTBF)

Annualized rate of occurrence (ARO)

Explanation:

RTO defines the maximum time an organization can tolerate downtime of a particular business process or system before the impact becomes unacceptable.

MTTR refers to the average time a device will take to recover from failure, not a concept directly resulting from a business impact analysis.

MTBF is the expected amount of time a device or product will last in operation between failures and does not factor into the determination of acceptable downtime in a business impact analysis.

ARO is a risk management term referring to the expected frequency of a risk occurring in a year.

q_risk_types_tol_rto_02_secp8

A healthcare institution is conducting a risk assessment to evaluate the implications of potential system downtime.

The team wants to pinpoint the maximum allowable duration that a crucial healthcare system can be nonfunctional before seriously affecting hospital operations.

What is the institution trying to establish in this situation?

Answers:

***Recovery time objective (RTO)**

Mean time to repair (MTTR)

Mean time between failures (MTBF)

Annualized rate of occurrence (ARO)

Explanation:

RTO defines the maximum time an organization can tolerate downtime of a particular business process or system before the impact becomes unacceptable.

MTTR refers to the average time a device will take to recover from failure, not a concept directly resulting from a business impact analysis.

MTBF is the expected amount of time a device or product will last in operation between failures and does not factor into the determination of acceptable downtime in a business impact analysis.

ARO is a risk management term referring to the expected frequency of a risk occurring in a year.

q_anylz_risk_ad_hoc_secp8

A company decides to conduct a risk assessment only once due to a specific change in its infrastructure. This decision to conduct the risk assessment is not part of its regular risk management process.

What type of risk assessment does this situation describe?

Answers:

***Ad hoc**

Recurring

Continuous

Quantitative

Explanation:

Ad hoc risk assessments are those conducted in response to specific organizational events or changes that are not part of the regular risk management process.

A recurring risk assessment is a type of risk assessment conducted on a regular, predetermined schedule. In this case, the company conducted the assessment once.

Continuous risk assessment is an approach where the company conducts assessments on a continuous, ongoing basis. This does not apply to the scenario in the question. It describes a one-time risk assessment.

Quantitative risk assessment involves using numerical measurements to assess the probability and consequences of risks. The type of risk assessment described in the question does not support methodology used to conduct the assessment.

q_anylz_risk_ale_01_secp8

You have conducted a risk analysis to protect a key company asset. You identify the following values:

Asset value = 400

Exposure factor = 75

Annualized rate of occurrence (ARO) = .25

Countermeasure A has a cost of 320 and protects the asset for four years. Countermeasure B has an annual cost of 85. An insurance policy to protect the asset has an annual premium of 90.

What should you do?

Answers:

***Accept the risk or find another countermeasure.**

Implement countermeasure A.

Implement countermeasure B.

Purchase the insurance policy.

Explanation:

In this scenario, you should either accept the risk or find a cheaper countermeasure. The cost of either countermeasure or the insurance policy exceeds the annualized loss expectancy (ALE) of the asset. The ALE = the asset value (400) x the exposure factor (.75) x the ARO (.25) = 75. Based on these calculations, you would expect an annual loss of 75.

Countermeasure A has an annual cost of 80, countermeasure B has an annual cost of 85, and the insurance premium has an annual cost of 90 (all of which cost more than the expected annual loss).

q_anylz_risk_ale_02_secp8

You have conducted a risk analysis to protect a key company asset. You identify the following values:

Asset value = 400

Exposure factor = 75

Annualized rate of occurrence = .25

What is the annualized loss expectancy (ALE)?

Answers:

25

***75**

100

175

475

Explanation:

To calculate the ALE, use the following formula:

Asset value (AV) x exposure factor (EF) x annualized rate of occurrence (ARO) => 400 x 75% x .25 = 75

q_anylz_risk_ale_03_secp8

A company's risk management team has identified a particular risk that carries a significant financial cost. The team has also determined the frequency at which this risk event is likely to occur over a year.

Based on these criteria, what is the company trying to calculate?

Answers:

***Annualized loss expectancy (ALE)**

Single loss expectancy (SLE)

Risk threshold

Exposure factor

Explanation:

The company calculates ALE by multiplying the SLE, the expected monetary loss from a single risk event, by the annualized rate of occurrence (ARO).

While the SLE is part of the ALE calculation, it represents the expected monetary loss from a single risk event and does not include the frequency of the event.

The risk threshold refers to the level of risk an organization is willing to accept or tolerate. It does not utilize the financial cost of a risk event and its frequency of occurrence in its calculations.

Exposure factor is a percentage that represents the magnitude of loss a potential risk event would have on a specific asset, not the annualized financial cost of a risk.

q_anylz_risk_ale_04_secp8

An organization's risk assessment team has flagged a particular risk event that poses a substantial financial threat in every financial quarter.

What is the team aiming to compute based on these parameters?

Answers:

***Annualized loss expectancy (ALE)**

Single loss expectancy (SLE)

Risk threshold

Exposure factor

Explanation:

To calculate ALE, the team multiplies the SLE, the expected monetary loss from a single risk event, by the annualized rate of occurrence (ARO).

While the SLE is part of the ALE calculation, it represents the expected monetary loss from a single risk event and does not include the frequency of the event.

The risk threshold refers to the level of risk an organization is willing to accept or tolerate. It does not utilize the financial cost of a risk event and its frequency of occurrence in its calculations.

Exposure factor is a percentage that represents the magnitude of loss a potential risk event would have on a specific asset, not the annualized financial cost of a risk.

q_anylz_risk_appetite_levels_secp8

As the chief risk officer of a rapidly growing tech startup, you are tasked with setting the company's risk appetite.

The company is in a strong financial position and is looking to aggressively expand its market share. The board of directors has expressed their willingness to take on more risk to achieve this goal.

Which of the following levels of risk appetite would you recommend?

Answers:

Minimalist

Cautious

Balanced

***Expansionary**

Aggressive

Explanation:

Expansionary is the correct answer. This level of risk appetite is correct because it involves a high tolerance for risk. It is suitable for organizations that are in a strong financial position and are looking to aggressively expand their market share, which aligns with the tech startup's situation in this scenario.

The minimalist level of risk appetite is incorrect because it involves taking on the least amount of risk possible. This approach is typically used by organizations that prioritize stability over growth, which is not the case for the tech startup in this scenario.

The cautious level of risk appetite is incorrect because it involves a low tolerance for risk. While it may be suitable for organizations that are focused on steady, incremental growth, it is not appropriate for a company that is looking to aggressively expand its market share.

The balanced level of risk appetite is incorrect because it involves a moderate level of risk. While it may be suitable for organizations that are looking to balance growth with stability, it is not appropriate for a company that is looking to aggressively expand its market share.

The aggressive level of risk appetite is incorrect because it involves the highest level of risk tolerance. While it may be suitable for organizations that are willing to take on significant risk for the potential of high returns, it could also lead to substantial losses. The board of directors has expressed their willingness to take on more risk, but they have not indicated a willingness to take on the highest level of risk.

q_anylz_risk_aro_01_secp8

What is the term for the average number of times that a specific risk is likely to be realized in a single year?

Answers:

Annualized loss expectancy

***Annualized rate of occurrence**

Estimated maximum downtime

Exposure factor

Explanation:

Annualized rate of occurrence (ARO) is the average number of times that a specific risk is likely to be realized in a single year.

Annualized loss expectancy (ALE) is ARO x SLE (single loss expectancy), which is the estimated per-year loss due to exposures.

Estimated maximum downtime sounds similar to maximum tolerable downtime or recovery time objective, neither of which are related to the average number of times a risk is likely to be realized.

Exposure factor is the percentage of value loss that is experienced due to an exposure rather than the number of times of exposure.

q_anylz_risk_aro_02_secp8

When conducting a risk assessment, how is the annualized rate of occurrence (ARO) calculated?

Answers:

Multiply the value of the asset by an exposure factor (EF).

***Through historical data provided by insurance companies and crime statistics.**

Divide the value of the asset by an exposure factor (EF).

Multiply the single loss expectancy (SLE) times the annual rate of occurrence (ARO).

Explanation:

The annualized rate of occurrence (ARO) is the likelihood of a risk occurring within one year. Historical data provides the basis for the statistical probability of the risk occurring. This information is frequently obtained from insurance companies, law enforcement agencies, and computer incident-monitoring organizations. ARO is typically expressed in percent or decimal form.

Single loss expectancy (SLE) is calculated by multiplying the value of the asset by an exposure factor (EF).

Dividing the value of the asset by an exposure factor (EF) is not a valid quantitative analysis formula.

Exposure factor (EF) is calculated by multiplying the single loss expectancy (SLE) times the annual rate of occurrence (ARO).

q_anylz_risk_aro_03_secp8

A company's risk management team has been analyzing a potential risk to its operations. They have identified the probability of the risk event occurring, and they wish to express this probability on a yearly basis.

What is the company trying to calculate?

Answers:

***Annualized rate of occurrence (ARO)**

Annualized loss expectancy (ALE)

Single loss expectancy (SLE)

Risk threshold

Explanation:

The ARO represents the estimated expected frequency that a risk occurs on a yearly basis.

The ALE calculation combines the ARO with the SLE, representing the annual expected loss from a risk, but is not a direct translation of a risk's probability to a yearly rate.

SLE represents the financial impact of a risk event occurring once, not the frequency or probability of the risk over a year.

The risk threshold is the level of risk an organization is willing to accept, which is different from the concept of estimating the frequency of a risk event on a yearly basis.

q_anylz_risk_aro_04_secp8

A financial institution's risk assessment department evaluates a potential threat that could disrupt its financial transactions. They have quantified the likelihood of this risk event happening and aim to project this probability over an annual period.

What metric are they attempting to establish?

Answers:

***Annualized rate of occurrence (ARO)**

Annualized loss expectancy (ALE)

Single loss expectancy (SLE)

Risk threshold

Explanation:

The ARO represents the estimated expected frequency that a risk occurs on a yearly basis.

The ALE calculation combines the ARO with the SLE, representing the annual expected loss from a risk, but is not a direct translation of a risk's probability to a yearly rate.

SLE represents the financial impact of a risk event occurring once, not the frequency or probability of the risk over a year.

The risk threshold is the level of risk an organization is willing to accept, which is different from the concept of estimating the frequency of a risk event on a yearly basis.

q_anylz_risk_calculate_secp8

Which of the following are main variables in calculating risks? (Select two.)

Answers:

***Probability**

***Likelihood**

Risk mitigation cost

Risk response time

Risk assessment

Explanation:

The following are main variables in calculating risks:

Probability is a quantitative measure typically expressed as a numerical value between 0 and 1 or a percentage. Probability aims to precisely measure the chance of a risk event occurring based on statistical methods. Therefore, it is a main variable in calculating risks.

Likelihood is often used in qualitative analysis to describe the chance of a risk event happening subjectively. It is typically expressed using "low," "medium," and "high" or scored on a scale from 1 to 5. Therefore, it is also a main variable in calculating risks.

Risk mitigation cost is a factor considered after risks have been identified and assessed. It is not a variable used in the calculation of risk itself.

Risk response time is a factor that comes into play after a risk has occurred. It is not a variable used in the calculation of risk itself.

Risk assessment is a core component of a cybersecurity program that evaluates previously identified risks to determine their potential impact on the organization. Calculating risks is a part of that component.

q_anylz_risk_ef_01_secp8

A cybersecurity team is conducting a risk assessment of a company's IT infrastructure. The team identified a potential vulnerability that could result in the loss of a critical system.

They determined if the team does not remediate the issue, they would lose approximately 40 percent of the system's data.

What is this percentage representative of in this scenario?

Answers:

***Exposure factor (EF)**

Single loss expectancy (SLE)

Annualized loss expectancy (ALE)

Impact

Explanation:

The EF is the percentage of loss that a realized threat would have on certain assets. The cybersecurity team determined a 40% expected data loss in the scenario.

SLE is the product of the asset value and the EF. The scenario mentions only the EF (40 percent of data loss), not the overall SLE.

Calculate the ALE by multiplying SLE by the annualized rate of occurrence (ARO). The scenario does not provide information to calculate ALE.

The term impact in a risk assessment context typically refers to the consequences or effects of a risk if it occurs, not the percentage of loss, which is the exposure factor.

q_anylz_risk_ef_02_secp8

A company wants to determine the single loss expectancy (SLE) for a critical server.

What formula will the company use to calculate the SLE?

Answers:

***Asset x exposure factor (EF)**

Asset x annualized rate of occurrence

Annualized loss expectancy (ALE) x exposure factor (EF)

Annualized loss expectancy (ALE) x annualized rate of occurrence

Explanation:

The SLE is the amount a company would lose in a single risk factor occurrence. The calculation for SLE is the multiplication of the asset's value by an exposure factor (EF).

Annualized loss expectancy (ALE) is the amount a company would lose over the course of a year. The calculation for ALE is the multiplication of SLE by the annualized rate of occurrence (ARO).

To determine the ALE, the company will need to determine the SLE first. The ALE multiplied by the EF cannot determine the SLE.

The ARO calculates the ALE and cannot calculate the SLE.

q_anylz_risk_heat_map_01_secp8

A company plans to upgrade its wireless network infrastructure to improve connectivity and security. The IT team wants to ensure that the new network design provides adequate coverage, minimizes interference, and meets security standards.

To achieve this, they conduct a site survey and create a heat map of the area.

What is the primary purpose of conducting a site survey and creating a heat map for the company's wireless network upgrade?

Answers:

***To assess wireless signal coverage, identify dead zones, and optimize access point placement for the new network design.**

To identify potential security threats and vulnerabilities in the existing network.

To map out the physical layout of the building and identify potential obstacles that could affect wireless signal strength.

To evaluate the performance and bandwidth usage of the current wireless network.

Explanation:

Conducting a site survey and creating a heat map allows the IT team to assess the wireless network's coverage, identify dead zones, and optimize the placement of access points for the new network design. This process helps ensure that the new wireless infrastructure provides adequate coverage and minimizes potential connectivity issues.

The primary focus of the site survey and heat map is not on identifying security threats but on optimizing wireless signal coverage.

While mapping the physical layout is part of the process, the main purpose is to optimize coverage.

The survey is not primarily about evaluating performance and bandwidth usage but about coverage and placement optimization.

q_anylz_risk_heat_map_02_secp8

An agency needs to enhance its wireless network infrastructure to improve its overall security and connectivity. As part of the preparation, the IT Team decides to perform a site survey and develop a heat map of the site.

What is the main objective behind conducting the site survey and creating the heat map for the company's wireless network upgrade?

Answers:

***To assess wireless signal coverage, identify dead zones, and optimize access point placement for the new network design.**

To identify potential security threats and vulnerabilities in the existing network.

To map out the physical layout of the building and identify potential obstacles that could affect wireless signal strength.

To evaluate the performance and bandwidth usage of the current wireless network.

Explanation:

Conducting a site survey and creating a heat map allows the IT team to assess the wireless network's coverage, identify dead zones, and optimize the placement of access points for the new network design. This process helps ensure that the new wireless infrastructure provides adequate coverage and minimizes potential connectivity issues.

The primary focus of the site survey and heat map is not on identifying security threats but on optimizing wireless signal coverage.

While mapping the physical layout is part of the process, the main purpose is to optimize coverage.

The survey is not primarily about evaluating performance and bandwidth usage but about coverage and placement optimization.

q_anylz_risk_kri_01_secp8

A risk manager at a large corporation conducts a risk and control self-assessment (RCSA) to identify and assess risks for a new market expansion.

The manager identifies risks associated with new regulatory requirements, market volatility, and aggressive competitors.

Which key risk indicator (KRI) metric is MOST critical for managing these risks?

Answers:

The number of new customers acquired in the new market.

***The number of regulatory violations reported by the regulatory authority.**

The level of market volatility.

The frequency of security incidents involving sensitive customer data.

Explanation:

The most critical metric is the number of reported regulatory violations. The risk manager must prioritize ensuring regulatory compliance in a new market. Non-compliance can result in fines, damage to reputation, and operational disruption.

While important for business growth, customer acquisition does not directly measure risk exposure.

The company cannot directly control market volatility, an external factor. Instead, the company should develop strategies to manage the effects of volatility.

Security incidents pose significant risks, but regulatory compliance is the more immediate concern during expansion. The risk manager should manage security risks, but in this context, the focus should be on regulatory risks.

q_anylz_risk_kri_02_secp8

A company conducts a risk analysis on a newly developed software application and has identified various potential threats and their impacts.

Now, they want to set some measurable metrics that will help them understand when they might be approaching an unacceptable level of risk.

What is the company looking to define in this situation?

Answers:

***Key risk indicators (KRIs)**

Risk threshold

Risk tolerance

Risk appetite

Explanation:

Organizations use KRIs as measurable metrics to provide an early signal of increasing risk exposures in various areas of the enterprise.

While risk thresholds involve levels of risk that are acceptable to an organization, they do not involve using measurable metrics that provide an early warning of risk.

Risk tolerance is an important concept in risk management, although it doesn't involve defining measurable metrics to understand risk levels.

Risk appetite is the acceptable level of risk an organization will accept in pursuit of its objectives before action is necessary to reduce the risk. However, it does not utilize measurable metrics or indicators.

q_anylz_risk_quantitative_01_secp8

When analyzing assets, which analysis method assigns financial values to assets?

Answers:

Qualitative

***Quantitative**

Transfer

Acceptance

Explanation:

Quantitative analysis assigns a financial value or assignment of real numbers and the cost required to recover from a loss to the asset.

Qualitative analysis seeks to identify costs that cannot be concretely defined using quantitative analysis.

Transfer and acceptance are responses to risk; they are not risk analysis methods.

q_anylz_risk_quantitative_02_secp8

Which of the following does a quantitative risk assessment include? (Select the three best options.)

Answers:

***Single loss expectancy (SLE)**

***Annual rate of occurrence**

***Annual loss expectancy (ALE)**

Recovery time objective

Heat map

Risk register

Key risk indicators (KRIs)

Explanation:

A quantitative risk assessment includes:

The single loss expectancy (SLE) representing the cost of any single item loss. It is both its own formula and used to calculate ALE.

The annual rate of occurrence (ARO) indicating the number of times the loss will occur within a year. It is both its own formula and used to calculate ALE.

The annual loss expectancy (ALE) which is the single loss expectancy (SLE) times the annual rate of occurrence (ARO). ALE is the main purpose of the quantitative risk assessment.

The following are not included in a quantitative risk assessment:

The recovery time objective (RTO) which identifies the maximum time it takes to recover a system when an outage occurs.

A heat map is a simple approach to measuring inherent risk. For each risk, a simple red, yellow, or green indicator can be put into each column to represent the severity of the risk, its likelihood, cost of controls, and so on.

A risk register is a document showing the results of risk assessments in a comprehensible format and includes information regarding risks, their severity, the associated owner of the risk, and all identified mitigation strategies.

Key risk indicators (KRIs) are critical predictive indicators organizations use to monitor and predict potential risks.

q_anylz_risk_register_01_secp8

After reading an article online, a concerned business stakeholder wishes to discuss the risk associated with denial-of-service (DoS) attacks.

The stakeholder requests information about the possibilities of an attacker learning about the countermeasures in place.

Where would the security analyst look to find this information?

Answers:

***Risk register**

Risk heat map

Risk regulations

Risk and control assessment (RCA)

Explanation:

The risk register shows risk assessment results in a comprehensible format. Information in the register includes impact, likelihood ratings, date of identification, description, countermeasures, owner/route for escalation, and status.

The risk heat map is a graphical table indicating the likelihood and impact of risk factors identified. The heat map alone does not include countermeasures.

Regulations affect risk posture but do not include specifics related to specific risks or a company. Regulations include requirements to deploy security controls and make demonstrable efforts to reduce risk.

An RCA is a process that identifies risks and assesses control effectiveness intended to mitigate those risks. This process does not define countermeasures used when a company finds a risk.

q_anylz_risk_register_02_secp8

A technology firm is adopting a structured approach to managing risk. It catalogs potential risks and assigns them to responsible parties within the team. This catalog also includes the strategies to mitigate these risks and their potential impacts.

What document must the firm develop as it adopts this approach?

Answers:

***Risk register**

Risk assessment

Risk analysis

Business impact analysis

Explanation:

A risk register is a document that captures identified risks and includes their definition, the effect of their occurrence, the response strategy, and a risk owner.

Risk assessment does not inherently involve assigning responsibility for those risks to individuals or tracking them over time, a function of the risk register.

Risk analysis involves examining how project outcomes and objectives might change due to the impact of the risk event. It does not include the assignment of risk owners or the documentation of risks.

Business impact analysis involves identifying potential impacts and their effects; it does not inherently involve assigning responsibility for those impacts to individuals or tracking them over time.

q_anylz_risk_register_03_secp8

An information security officer creates a document that identifies downtime, the likelihood of occurrence, the probable impact of the downtime, and steps to mitigate the downtime and scores the likelihood based on defined security controls.

What is the information security officer creating in this instance?

Answers:

***Risk register**

Supply chain assessment

Vulnerability assessment

Supply chain analysis

Explanation:

A risk register is a repository for documenting risks identified in an organization and includes information and steps to take regarding the risk. Common information found in a risk register is the specific risk, the likelihood of occurrence, and the action.

A supply chain assessment evaluates all elements required to produce and distribute a product. The goal of a supply chain assessment is to identify areas where a company can improve.

A vulnerability assessment identifies and provides mitigation for system vulnerabilities in a company.

Supply chain analysis evaluates the risks and vulnerabilities associated with the various entities (vendors) involved in a supply chain by examining the security practices, capabilities, and reliability of individual vendors within the supply chain network.

q_anylz_risk_reporting_secp8

What is the primary purpose of risk reporting in an organization?

Answers:

***To inform stakeholders about the current risk status.**

To assign blame for risks that have occurred.

To predict future risks with 100% accuracy.

To eliminate all risks in the organization.

Explanation:

The primary purpose of risk reporting is to inform stakeholders about the current risk status. It provides an overview of the identified risks, their potential impact, and the measures taken to mitigate them. This allows stakeholders to make informed decisions about risk management.

Assigning blame for risks that have occurred is not the purpose of risk reporting. Risk reporting is a proactive process aimed at managing risks, not a reactive process for blaming individuals or teams when risks occur.

While risk reporting can help in predicting potential future risks, it cannot do so with 100% accuracy. Risk prediction involves a degree of uncertainty and the goal is to estimate potential risks based on available data, not to predict them with absolute certainty.

The goal of risk reporting is not to eliminate all risks in the organization, as this is practically impossible. Instead, it aims to manage risks effectively by identifying them, assessing their potential impact, and taking appropriate measures to mitigate them.

q_anylz_risk_sle_01_secp8

Which of the following BEST defines single loss expectancy (SLE)?

Answers:

The statistical probability of a malicious event.

The monetary value of a single employee's loss of productivity due to a successful attack.

***The total monetary loss associated with a single occurrence of a threat.**

The total cost of all countermeasures associated with protecting against a given vulnerability.

Explanation:

Single loss expectancy (SLE) is best defined as the total monetary loss associated with a single occurrence of a threat. The key to this definition is the term total. In other words, this encompasses all costs, including lost employee productivity, replacement hardware/software, and payroll for additional consultants. All of this must be considered when calculating the total loss.

Probability is a quantitative measure typically expressed as a numerical value between 0 and 1 or a percentage. Probability aims to precisely measure the statistical probability of a malicious event occurring.

The monetary value of a single employee's loss of productivity due to a successful attack is simply a statistic that is often included in risk reporting.

Inherent risk includes accounting for security controls or countermeasures that could be introduced to address every risk factor.

q_anylz_risk_sle_02_secp8

You have conducted a risk analysis to protect a key company asset. You identify the following values:

Asset value = 400

Exposure factor = 75

Annualized rate of occurrence = .25

What is the single loss expectancy (SLE)?

Answers:

100

***300**

475

30000

Explanation:

The single loss expectancy (SLE) is the asset value (AV) multiplied by the exposure factor (EF), with the EF being a percentage of the asset value that is lost. In this example, SLE = 400 x 75% = 300.

q_anylz_risk_sle_03_secp8

A tech start-up company considers deploying a new email system. The start-up is currently identifying risks associated with the potential downtime of the new system and considers the costs for each event.

What metric should the company utilize during this process?

Answers:

***Single loss expectancy (SLE)**

Annualized rate of occurrence (ARO)

Annualized loss expectancy (ALE)

Risk identification

Explanation:

Single loss expectancy (SLE) is a monetary value assigned to a single event representing the company's potential loss amount if a specific threat occurred.

The annualized rate of occurrence (ARO) refers to the estimated frequency where an expected threat will occur within a year. It does not reflect a monetary loss from a single event.

Annualized loss expectancy (ALE) is the potential yearly cost of all instances of a specific threat, calculated by multiplying the SLE with the ARO.

While risk identification is a key aspect of risk management and is the process of finding, recognizing, and describing risks, it does not refer to estimating the financial loss from a single event.

q_anylz_risk_sle_04_secp8

A newly established software firm plans to roll out a proprietary messaging platform. One of their main concerns is estimating the platform's monetary impact per incident of downtime.

What measure should the firm employ to address this concern?

Answers:

***Single loss expectancy (SLE)**

Annualized rate of occurrence (ARO)

Annualized loss expectancy (ALE)

Risk identification

Explanation:

SLE is a monetary value assigned to a single event that represents the company's potential loss amount if a specific threat occurs.

ARO refers to the estimated frequency at which a threat occurs within a year. It does not reflect a monetary loss from a single event.

ALE is the potential yearly cost of all instances of a specific threat, calculated by multiplying the SLE with the ARO.

While risk identification is an important aspect of risk management and is the process of finding, recognizing, and describing risks, it does not refer to the estimated financial loss from a single event.

q_anylz_risk_sle_ale_secp8

A tech company employs the single loss expectancy (SLE) and annualized loss expectancy (ALE) models for quantitative assessment and uses subjective judgment for qualitative analysis.

They use a heat map or traffic light impact matrix to represent the severity of the risk, its likelihood, cost of controls, etc.

What is the primary benefit of the company's approach of combining both quantitative and qualitative risk assessment methods?

Answers:

It allows for a quick initial assessment of risks and focuses on the most significant issues.

It develops tangible numbers that reflect real money and justifies the costs of various controls.

***It provides both numerical data for precision and subjective judgment for situations in which precise data is unavailable.**

It eliminates the subjectivity in risk assessment completely.

Explanation:

The company's approach employs both numerical data for precision (quantitative) and subjective judgment for situations in which precise data is unavailable (qualitative). This mixed approach provides a comprehensive understanding of the risks, their potential impact, and the likelihood of their occurrence.

While this statement accurately describes a benefit of qualitative risk analysis, it does not fully capture the advantage of combining both quantitative and qualitative methods.

This statement describes the benefit of a quantitative risk analysis approach. The combined approach of using both methods provides broader benefits.

Even though quantitative risk analysis uses numerical data for precision, the combined approach does not eliminate subjectivity completely.

q_anylz_risk_threshold_secp8

A company determines a certain level of risk that, once exceeded, requires immediate action or reconsideration of the initiative. The company takes pride in its cautious approach to business and generally avoids high-risk activities.

Which of the following should the company employ to align with its desired risk management approach?

Answers:

***Risk threshold**

Risk tolerance

Risk appetite

Risk mitigation

Explanation:

The risk threshold refers to a specific point at which risk becomes unacceptable, necessitates intervention, and the company requires immediate action or reconsideration of the initiative.

Risk tolerance refers to the degree of variability in investment returns an organization is willing to withstand. It doesn't necessarily correspond to a predetermined required point for action.

Risk appetite refers to the total exposed amount an organization wishes to undertake based on risk-return trade-offs for its business objectives.

Risk mitigation does not capture the concept of setting a defined level (threshold) beyond which risk is unacceptable.

q_biz_cont_backup_secp8

A healthcare provider is digitalizing its patient data and needs a system that guarantees the preservation of its sensitive data in the event of a disaster while avoiding lock-in to a single technology vendor's platform.

What is the MOST crucial practice the provider needs to implement?

Answers:

***Maintaining regular data backups**

Implementing an uninterruptible power supply (UPS)

Adopting platform diversity for their digital infrastructure

Incorporating additional power generators

Explanation:

Regular data backups provide a reliable way to restore data if the original becomes compromised, lost, or corrupted due to a disaster.

A UPS system ensures continuity of operations in the face of smaller-scale disruptions, like power outages, rather than handling large-scale data loss scenarios.

Platform diversity is a strategic decision about technology procurement and management rather than a specific data preservation measure.

Incorporating additional power generators maintains service availability during power disruptions. They do not preserve data integrity during a disaster or contribute to avoiding vendor lock-in.

q_biz_cont_bcp_01_secp8

What is the primary goal of business continuity planning (BCP)?

Answers:

Minimize the organization's risk of service delays and interruptions.

***Maintain business operations with reduced or restricted infrastructure capabilities or resources.**

Minimize decision-making during the development process.

Protect an organization from major computer services failure.

Explanation:

The primary goal of BCP is to maintain business operations with reduced or restricted infrastructure capabilities or resources.

Minimizing the risk of service delays and interruptions is a goal of DRP, disaster recovery plan. If your organization cannot provide services, it is experiencing a disaster. Minimizing decision making during the development process is not a valid goal of BCP or DRP; decisions should be made during development.

The correct DRP goal is to minimize decisions during an emergency.

Protecting an organization from major computer services failure is a goal of DRP, not BCP. If computer services fail, business continuity is interrupted, creating a disaster.

q_biz_cont_bcp_02_secp8

When is a BCP or DRP design and development actually completed?

Answers:

Once senior management approves

***Never**

Only after testing and drilling

Only after implementation and distribution

Explanation:

BCP (business continuity plan) and DRP (disaster recovery plan) developments are never complete as they need constant improvement and updates.

Senior management approval, testing, drilling, implementation, and distribution are all important phases and elements in the life of BCP and DRP projects. However, they do not represent the end point of BCP/DRP design and development.

q_biz_cont_bcp_03_secp8

In a recent cybersecurity meeting, the IT director emphasized prioritizing business continuity planning (BCP). The organization seeks to ensure uninterrupted business processes, especially in the wake of potential disruptions or disasters.

Which primary objective BEST captures the essence of BCP in an organization's overall security strategy?

Answers:

***Ensure that critical business processes remain operational during and after disruptions.**

Establish secure connections between office locations.

Detect and respond to any unauthorized system access.

Encrypt sensitive company data stored in databases.

Explanation:

BCP maintains the continuity of essential business functions in the face of disruptions, whether due to natural disasters, cybersecurity incidents, or other unforeseen challenges.

Although secure connections are essential for data protection and integrity, they do not represent the primary goal of BCP. This aspect is more related to secure networking and remote access policies.

Detection and response to unauthorized access are vital components of security, but they are primarily the roles of intrusion detection and response systems, not BCP.

Encryption ensures the confidentiality of data, but it is not the main focus of BCP. Data encryption is more about data at rest protection and not about business continuity.

q_biz_cont_bcp_04_secp8

A large business works with a consulting group to develop a business continuity plan. The goal of the plan is to provide a potentially uninterrupted workflow in the event of an incident.

Which one of the descriptions matches what the business is trying to achieve?

Answers:

***Ensuring processing redundancy supports the workflow**

Performing mission-critical functions without IT support

Recovery of secondary business functions when disrupted

Retention of data for a specified period

Explanation:

Business continuity planning identifies how business processes should deal with both minor and disaster-level disruption. It ensures that there is processing redundancy supporting the workflow through failover.

Organizations use continuity of operation planning (COOP) in government facilities. In some definitions, COOP refers specifically to backup methods of performing mission functions without IT support.

A disaster recovery plan is necessary when an organization's primary business function goes down. Disaster recovery requires considerable resources, such as shifting processing to a secondary site.

A retention policy is important for retrospective incident handling. A company uses the retention policy for historic logs and for how long to keep the data.

q_biz_cont_bcp_coop_secp8

During a cybersecurity seminar, the IT manager presented two significant components of their organization's continuity strategy: business continuity planning (BCP) and continuity of operations planning (COOP). The team needed clarification about the distinctive attributes of each component.

Which statement BEST distinguishes business continuity planning from continuity of operations planning in the context of an organization's overall continuity approach?

Answers:

***BCP focuses on the recovery and continuity of business functions, while COOP emphasizes maintaining essential operations during a disruption.**

BCP deals exclusively with data backups, while COOP handles physical infrastructure.

COOP is solely concerned with natural disasters, while BCP addresses cyberattacks.

BCP ensures a company's long-term profitability, while COOP only addresses short-term operational goals.

Explanation:

BCP serves to recover and ensure the continuity of business functions after a disruption. In contrast, COOP ensures that essential operations remain during the disruption itself.

Both BCP and COOP may encompass data backups and physical infrastructure. They do not exclusively handle one or the other.

Both COOP and BCP address a range of threats, including natural disasters and cyberattacks. They are not exclusive to one type of threat.

While BCP does consider long-term recovery and profitability, COOP focuses on immediate operational continuity, not just short-term goals. Both are crucial for an organization's sustainability.

q_biz_cont_capacity_planning_secp8

A prominent e-commerce company experiencing significant business growth anticipates a sharp increase in website traffic during an upcoming annual sales event. The company is wary of potential system bottlenecks or downtimes that could disrupt sales and affect reputation.

What primary strategy should the company use to ensure its systems can handle the upcoming event?

Answers:

***Rigorous capacity planning process**

Installing a uninterruptible power supply (UPS)

Implementing a failover system

Expansion of the infrastructure at the primary data center

Explanation:

The goal of capacity planning is to ensure that the system can cope with growth in demand without exceeding its capacity.

Uninterruptible power supplies (UPS) do not directly impact system performance or capacity to handle increased loads.

Failover systems ensure that services remain available during a system failure but do not contribute to the system's overall processing power or bandwidth.

Expanding the primary data center infrastructure might be part of the solution, but it is not a strategic approach to managing increased system loads, especially if the company only requires expansion at certain operation times. This would otherwise be costly.

q_biz_cont_cloud_secp8

You are the IT manager of a mid-sized company that has recently migrated its data and applications to the cloud.

A major natural disaster has just occurred, causing significant damage to your company's physical infrastructure. You need to ensure business continuity and minimize downtime.

What is the MOST appropriate action to take?

Answers:

Immediately purchase new hardware to replace the damaged infrastructure.

Rely on the cloud service provider to handle all aspects of disaster recovery.

***Initiate the disaster recovery plan that includes the use of cloud services.**

Wait for the physical infrastructure to be repaired before resuming operations.

Explanation:

Initiating the disaster recovery plan that includes the use of cloud services is the correct answer as it would ensure business continuity and minimize downtime. The cloud allows for data and applications to be accessed from anywhere, which is particularly useful in a disaster scenario where physical infrastructure may be damaged.

Immediately purchasing new hardware to replace the damaged infrastructure is not the most appropriate because the company has already migrated its data and applications to the cloud. Purchasing new hardware would be unnecessary and costly.

While the cloud service provider plays a significant role in disaster recovery, it's not advisable to rely solely on them. The company should have its own disaster recovery plan in place.

Waiting for the physical infrastructure to be repaired before resuming operations would result in significant downtime and potential loss of business. With the data and applications already in the cloud, operations can continue despite the damage to the physical infrastructure.

q_biz_cont_documentation_secp8

As the business continuity manager of a large corporation, you are reviewing the company's business continuity plan (BCP).

Which of the following elements is NOT typically included in a comprehensive BCP?

Answers:

A list of critical business functions and the resources needed to support them.

***The company's annual financial statements.**

Detailed recovery strategies and procedures.

Contact information for key personnel and stakeholders.

Explanation:

While annual financial statements are important documents within a company, they are not typically part of a business continuity plan. The BCP focuses on how to maintain and restore business operations in the event of a disruption, not on the company's financial performance.

A list of critical business functions and the resources needed to support them is a fundamental part of a BCP. It helps to prioritize recovery efforts in the event of a disruption.

Detailed recovery strategies and procedures are a crucial part of a BCP. They provide a roadmap for how to restore operations following a disruption.

Contact information for key personnel and stakeholders is typically included in a BCP. This ensures that all relevant parties can be contacted quickly in the event of a disruption.

q_biz_cont_failover_secp8

As part of the robust disaster recovery plan, a major e-commerce company is ensuring that in the event of a server failure, its customer transactions and data remain unaffected, and operations continue without any noticeable interruption.

Which of the following strategies should the company prioritize?

Answers:

***Implementing a failover mechanism**

Conducting regular tabletop exercises

Expanding the capacity of its uninterruptible power supply (UPS) systems

Increasing the frequency of its data backups

Explanation:

A failover mechanism allows for an automatic switch to a standby system upon the failure of the previously active server, ensuring uninterrupted service.

Tabletop exercises affect disaster recovery techniques but do not physically prevent service interruption.

While an uninterruptible power supply (UPS) provides emergency power to a load when the input power source fails, it does not ensure the continuity of service in the event of server failure.

Regular data backups are essential for data recovery, but they will not ensure continuous operation in the event of a server failure.

q_biz_cont_growth_secp8

In a recent company meeting, the network administrator discussed the upcoming growth projections for the next year. The IT team ensures that the organization's infrastructure can sustain the anticipated growth in user demand and traffic volume.

Which risk poses the MOST significant threat if the IT team fails to address the infrastructure needs for the forecasted growth adequately?

Answers:

An increase in phishing attack susceptibility

Reduced efficiency of network address translation (NAT)

Lesser demand for cloud storage services

***Potential for service degradation or unavailability**

Explanation:

Without properly addressed capacity planning, the primary risk involves not being able to handle the increased demand, leading to service degradation or even complete unavailability.

While security measures are crucial, capacity planning primarily focuses on scalability and performance. Phishing susceptibility is not directly related to capacity planning.

NAT efficiency is more related to routing and addressing. While capacity can impact every aspect of IT, NAT is not the primary concern related to capacity planning.

The demand for cloud storage services focuses on data needs rather than directly tied to capacity planning related to user demand and traffic volume.

q_biz_cont_hot_site_secp8

A company wants to guarantee business continuity even if a catastrophic event occurs at its primary data center.

What are the MOST appropriate strategies for the company to adopt? (Select two.)

Answers:

***Deploy a hot site**

***Disperse locations geographically**

Deploy a warm site

Utilize multicloud systems

Implement capacity planning

Explanation:

The following are the most appropriate strategies for the company to adopt:

A hot site is a real-time replica of the primary system, providing immediate failover capability. If the main system fails, the hot site instantly takes over, minimizing downtime and ensuring uninterrupted business operations.

Geographic dispersion offers protection against region-specific threats, mitigating risks and enhancing uninterrupted service.

A warm site is a more cost-effective disaster recovery option than a hot site but requires time to become operational after a failure.

Multi-cloud strategies use multiple cloud providers to enhance resilience and flexibility. However, they do not provide an immediate backup for physical data centers.

Capacity planning is a critical process in which organizations assess their current and future resource requirements to ensure they can efficiently meet their business objectives. However, it is not a business continuity factor when a catastrophic event occurs at its primary data center.

q_biz_cont_irp_secp8

During an incident response, a company discovers several problems with its incident response plan (IRP), including the communication plan being incomplete, information missing, and have assigned people that are no longer employed with the company.

What can the company do BEST to mitigate these problems from occurring? (Select the three best options.)

Answers:

***Conduct a tabletop exercise**

***Create a simulation**

***Initiate a walkthrough**

Review IRP after incidents

Review capacity planning

Improve continuity of operations

Add storage space for backups

Explanation:

The following are steps the company can take to mitigate these problems from occurring again:

A tabletop exercise involves gathering required personnel to talk through steps they would take when responding to incidents. After the completion of the exercise, the organization updates the IRP.

A simulation is an exercise where red team attempts an intrusion, blue team operates response and recovery controls, and white team moderates the exercise. They are an excellent way to test the IRP.

A walkthrough presents a scenario to responders to demonstrate actions they would complete. It involves scans and actions executed on a sandbox of the company's actual response tools.

Only reviewing the IRP after incidents will not allow the company to make changes when personnel move positions, or they add new software.

Capacity planning is a critical process in which organizations assess their current and future resource requirements to ensure they can efficiently meet their business objectives. However, it is not a part of incident response.

Continuity of operations (COOP) refers to the process of ensuring that an organization can maintain or quickly resume its critical functions in the event of a disruption, disaster, or crisis. While this process may help with an incident response, it is not a primary or best action to take.

While backups and storage space play a critical role in the continuity of operations by safeguarding against data loss and restoring systems and data in the event of disruptions, increasing storage space is not a best step to take to improve the IRP.

q_biz_cont_mission_secp8

When recovering from a disaster, which services should you stabilize first?

Answers:

Least business-critical services

***Mission-critical services**

Financial support

Outside communications

Explanation:

The services you should restore first are mission-critical services. If mission-critical services are not restored within their maximum tolerable downtime, the organization is no longer viable.

Least business-critical services should be restored last.

Financial support and outside communications should be restored only after all other services with a higher level of criticality are restored.

q_biz_cont_regulations_secp8

A multinational corporation operates in several countries with diverse regulations regarding data privacy and security.

What is the primary responsibility of the security team concerning the multitude of governmental and regulatory entities influencing the corporation's operations?

Answers:

***Ensuring compliance with all applicable regulations and laws.**

Lobbying governmental entities for favorable policies.

Avoiding any interaction with regulatory entities to maintain operational secrets.

Shaping internal policies independently from external regulations.

Explanation:

The security team's obligation is to ensure that the corporation complies with all relevant laws and regulations. Compliance is a key part of cybersecurity, particularly for multinational corporations, which can be subject to a multitude of regulations depending on their operational jurisdictions.

Lobbying governmental entities for favorable policies is typically a role for legal or government relations teams, not the security team.

Avoiding any interaction with regulatory entities to maintain operational secrets is incorrect because regulatory compliance often necessitates communication and cooperation with regulatory entities.

Shaping internal policies independently from external regulations is also incorrect because the security team must align internal policies with external regulations to maintain compliance and avoid penalties.

q_biz_cont_remote_work_secp8

In the context of a global manufacturing firm transitioning to a remote work arrangement due to a crisis, which aspect is the MOST critical to ensure business continuity?

Answers:

Identifying potential bottlenecks through regular trend analysis.

Maintaining the aesthetic value of their digital platforms.

Replacing employees unable to adapt to remote work.

***Developing robust remote work plans with appropriate technologies.**

Explanation:

Establishing robust remote work plans is crucial to ensure business continuity in a crisis when physical presence at work is impossible.

While regular trend analysis is an essential part of capacity planning to identify patterns and potential problems, the most critical step in a crisis would be to ensure that robust remote work plans are in place.

Maintaining the aesthetic value of its digital platforms, while important for user experience and branding, is not the most critical aspect in ensuring business continuity during a crisis.

Hiring new employees to replace those unable to adapt to remote work is not only time-consuming but also a costly endeavor that may not be feasible or practical in a crisis.

q_biz_cont_resilience_recovery_secp8

A lead architect is designing a new security system for a multinational corporation. The chief executive officer (CEO) emphasizes that the continuity of business operations is a top priority.

Why would incorporating resilience and recovery into the security architecture be vital in this scenario?

Answers:

It increases system efficiency.

It reduces system costs.

***It ensures system functionality during and after disruptions.**

It enhances system aesthetics.

Explanation:

Resilience implies that systems, applications, and services can recover quickly and continue operating even under adverse conditions like cyber-attacks or equipment failures. Recovery ensures the organization has strategies and measures to restore systems, applications, data, and services after a disruption.

While system performance and efficiency are important considerations, they do not directly relate to business continuity.

While resilience and recovery might indirectly help reduce costs by minimizing downtime and potential loss of data or productivity, cost reduction is not the primary function of these elements in a security architecture.

Aesthetic value is irrelevant to security architecture. The security architecture's design focuses on functionality, security, and resilience against potential threats and disruptions.

q_biz_cont_simulation_secp8

A multinational company worries that its IT department is getting complacent regarding cybersecurity. The company begins working with an outside company to create an incident in a sandbox environment to gauge the IT department's response to a strong attack.

This situation represents what type of testing scenario?

Answers:

***Simulation**

Tabletop exercise

Walkthrough

Communication plan

Explanation:

A simulation exercise takes extensive investment and preparation but presents the most real-world scenario to test personnel against. It is the best choice to test a response to an incident.

A tabletop exercise consists of select personnel discussing what their department would do in the event of an incident. It is not the best means to gauge readiness.

A walkthrough is slightly more than a tabletop exercise. This testing scenario allows the running of actual scans and demonstrates some of the capabilities, but it is not the best option.

A communication plan is part of the incident response plan and dictates who can release information during and after an incident.

q_biz_cont_tabletop_exercise_secp8

During an annual review, a health services company's leadership aims to scrutinize its disaster response and data recovery protocols. They focus on effectiveness, hidden weaknesses, and clarity of employee roles during a disaster.

Which course of action would BEST serve these objectives?

Answers:

***Organizing tabletop exercises**

Expanding the IT department

Increasing the frequency of data backups

Installing larger uninterruptible power supply (UPS) systems

Explanation:

Tabletop exercises effectively test disaster response protocols and individual roles in a low-risk, discussion-based environment. All while allowing the company to identify potential weaknesses and improve its plans accordingly.

While having more IT personnel may enhance some aspects of disaster response, it does not guarantee a comprehensive evaluation of the protocols or individual responsibilities.

Though regular data backups are crucial for data recovery, they will not provide an evaluation of the overall disaster response protocols or clarify individual roles during a crisis.

Uninterruptible power supply (UPS) systems are important for ensuring power continuity in a crisis, but they do not contribute to evaluating the disaster response plan or understanding individual roles.

12.2 Vendor Management

As you study this section, answer the following questions:

What are three types of third-party relationships?

How does onboarding with a third-party create security risk?

What security risks should be considered on a daily or ongoing basis?

Why is it important to reevaluate security risks when offboarding?

The key terms for this section include:

Term	Definition
Due diligence	A legal principal that responsible parties have used best practice or reasonable care and have not been negligent in discharging their duties.
Conflict of interest	When an individual or organization has investments or obligations that could compromise their ability to act objectively, impartially, or in the best interest of another party.
Questionnaires	In vendor management, a structured means of obtaining consistent information, enabling more effective risk analysis and comparison.
Rules of Engagement (RoE)	A definition of how a pen test will be executed and what constraints will be in place. This provides the pen tester with guidelines to consult as they conduct their tests so that they don't have to constantly ask management for permission to do something.
Memorandum of Understanding (MOU)	Usually a preliminary or exploratory agreement to express an intent to work together that is not legally binding and does not involve the exchange of money.
Nondisclosure Agreement (NDA)	An agreement that stipulates that entities will not share confidential information, knowledge, or materials with unauthorized third parties.
Memorandum of Agreement (MOA)	Legal document forming the basis for two parties to cooperate without a formal contract (a cooperative agreement). MOAs are often used by public bodies.
Business Partnership Agreement (BPA)	Agreement by two companies to work together closely, such as the partner agreements that large IT companies set up with resellers and solution providers.
Master Service Agreement (MSA)	A contract that establishes precedence and guidelines for any business documents that are executed between two parties.
Service-level Agreement (SLA)	An agreement that sets the service requirements and expectations between a consumer and a provider.
Statement of Work (SOW)/Work Order (WO)	A document that defines the expectations for a specific business arrangement.

This section helps you prepare for the following certification exam objectives:

Exam	Objective

CompTIA Security+ SY0-701	5.1 Summarize elements of effective security governance. Procedures Onboarding/offboarding 5.3 Explain the processes associated with third-party risk assessment and management. Vendor assessment Penetration testing Right-to-audit clause Evidence of internal audits Independent assessments Supply chain analysis Vendor selection Due diligence Conflict of interest Agreement types Service-level agreement (SLA) Memorandum of agreement (MOA) Memorandum of understanding (MOU) Master service agreement (MSA) Work order (WO)/statement of work (SOW) Non-disclosure agreement (NDA) Business partners agreement (BPA) Vendor monitoring Questionnaires Rules of engagement

12.2.1 Managing Third Parties (Lesson Video)

Transcript:

It's common for one organization to work directly with a third party in a vendor, supply chain, or business partner relationship. These relationships often require the both parties' information systems to connect and integrate. Doing this exposes both network to risks. You'll want to pay careful attention to the onboarding phase, the ongoing operations phase, and the offboarding phase of this third-party relationship.

Let's take a look at the onboarding phase first. During this phase of the relationship, you'll want to sit down with individuals from the third-party organization. Compare each of your organization's security policies and infrastructure. Are the security policies for each organization similar? If there are substantial differences, you should discuss them. Do both organizations have similar incident response procedures? If there are differences, how will each party handle incident response? Are the security controls similar? Are the audit policies similar? Is the security posture of each organization compatible enough to work together, or will the integration expose vulnerabilities in one or both organizations? Finding answers to these questions helps you to identify and evaluate differences and potential risks up front.

There will always be risks, so you'll need to work with your organization's leadership to determine if the risks of this potential relationship outweigh the benefits. If so, it's okay to walk away. But if the benefits outweigh the risks, you'll want to move forward with the relationship.

After that decision is made, you need to create an interoperability agreement. This agreement is a collection of documents that identifies the key responsibilities in the relationship. The documents will specify who is going to perform key tasks, how those tasks will be measured, and how decisions will be handled, and other such information.

The first interoperability agreement document you need is a Service Level Agreement, or SLA. An SLA states what each organization in the relationship will do and at what level. It may also define how disputes will be managed, provide warranties, specify disaster recovery procedures, and specify when the agreement will be terminated.

Second, we need a Memorandum of Understanding, or MOU. An MOU provides a brief summary of the relationship. It describes which party is responsible for performing specific tasks . Basically, it says who is going to do what, and when they'll do it.

Next, we'll need a Business Partnership Agreement, or BPA. A BPA explains the rules surrounding the finances of the partnership. This could include details on how to handle contributions, withdrawals, financial reporting, and the profit and loss distribution.

Another important document is the Interconnection Security Agreement, or ISA. The ISA documents how the information systems of both organizations will connect and how they're going to share data. The ISA is more technical in nature. The SLA and the MOU define the relationship, whereas the BPA and the ISA specify how the project work is going to take place.

When onboarding is complete, you can start conducting business according to the terms of the agreement. During the ongoing operations phase, you'll want to keep any eye on several things. First, regularly verify compliance with the interoperability agreement documents. Make sure that each organization is doing what it said it would do.

In addition, conduct periodic vulnerability assessments to verify that the shared network connections have not exposed or created some type of security weakness. You should also conduct regular security audits to ensure that organizations are following the security requirements contained in the interoperability agreement.

Communication between the organizations is critical. Share the vulnerability assessment and security audit findings with everyone involved in the relationship in order to maintain risk awareness.

When the operations phase is complete and you've reached the end of your agreement, you're ready to end the relationship. This process is called offboarding. Although this is an important phase, it's probably the most neglected. During offboarding, you want to ensure that all the security doors you opened to facilitate the relationship are closed. From a technical standpoint, you'll want to reset or disable any open or shared connections that you may have with the third party. This could include VPNs, firewalls, or routers and switches that were set up to share information from your network.

In addition, disable any domain trust relationships established between the organizations. Disable any user or group accounts used by third parties to access your organization's data. Finally, reset passwords used by the third party to access applications on your network.

Leaving any of these doors open could potentially expose the information on your network. During the offboarding process, you'll also want to revisit the NDA, and you'll want to put all of your end-of-business agreements into a written format.
That's it for this lesson. In this video, we talked about managing third parties. We looked at the onboarding phase and the various documents commonly used. We talked about the ongoing operations phase of the relationship, and we discussed considerations for the offboarding phase to ensure that access to your system ends with the relationship.

12.2.2 Managing Third Parties Facts

Third-party risk assessment involves several important processes integral to effective risk management practices. These processes include vendor due diligence, risk identification and assessment, ongoing monitoring, and incident response planning. Vendor due diligence involves evaluating and selecting vendors based on their security practices, financial stability, regulatory compliance, and reputation. Risk identification and assessment include identifying potential risks associated with vendor relationships and assessing their potential impact on the organization's operations, data, and reputation. Ongoing monitoring ensures vendors maintain security controls, adhere to contractual obligations, and promptly address identified risks or vulnerabilities.

These processes are critical in risk management practices as they help organizations identify, assess, and mitigate risks associated with third-party relationships. Organizations can proactively manage and reduce risks by implementing robust third-party risk assessment processes, protecting assets, maintaining regulatory compliance, and fostering a safe and secure operational environment.

This lesson covers the following topics:

- Vendor selection
- Conflict of interest
- Vendor assessment methods
- Rules of engagement
- Legal agreements

Vendor Selection

Vendor selection practices must systematically evaluate and assess potential vendors to minimize risks associated with outsourcing or procurement. It typically includes several steps, such as identifying risk criteria, conducting due diligence , and selecting vendors based on their risk profile. Risk management practices aim to identify and mitigate risks related to financial stability, operational reliability, data security, regulatory compliance, and reputation. The goal is to select vendors who align with the organization's risk tolerance and demonstrate the capability to manage risks effectively.

Third-Party Relationships

In the modern business world, it’s very common for one organization to work directly with a third party in one of the following ways:

Vendor - A vendor is a company that sells an organization's goods or provides supplies an organization needs.

Supply chain - Supply chain relationships are collaborative relationships in which companies work together to achieve their operational objectives.

Business partner - A business partnership is an agreement between parties to operate a business together and to share its profits.

Third-Party Vendor Assessment

A third-party vendor refers to an external person or organization that provides goods, services, or technology solutions to another organization but operates independently. Third-party vendors play a significant role in business operations by offering specialized expertise, products, and services that support or enable the organization's own capabilities. Third-party vendors can range from technology, software, and cloud service providers to suppliers and contractors and collectively represent an organization's supply chain. Third-party vendors bring efficiency, cost-effectiveness, expertise, and innovation to organizations but also introduce potential risks as they may have access to sensitive data, infrastructure, or critical processes. Proper vendor assessment and continuous monitoring ensure third-party vendors adhere to security standards and regulatory compliance and fulfill their obligations to safeguard business operations from potential vulnerabilities and disruptions.

Vendor assessment is a critical component of Governance, Risk, and Compliance (GRC) frameworks and plays a pivotal role in maintaining the security of IT and business operations. Vendor assessment includes carefully evaluating third-party vendor capabilities, practices, and security measures before engaging in business partnerships or entrusting them with sensitive data and critical services. The significance of vendor assessment stems from the fact that organizations increasingly rely on external vendors for various aspects of their operations, such as technology solutions, cloud services, supply chain management, and outsourcing. By thoroughly assessing vendors, businesses can ensure that their partners adhere to established security standards, comply with regulatory requirements, and mitigate potential risks effectively. Engaging with vendors with weak security practices or inadequate risk management measures can introduce significant vulnerabilities.

A study by Ponemon Institute and Bomgar identified the following statistics:

Companies allow 89 vendors to access their networks weekly, on average.

69% of organizations have experienced a data breach due to vendor security shortcomings.

65% of respondents say it's hard to manage cybersecurity risks associated with third-party vendors.

64% of respondents said their organization focuses more on cost than security when outsourcing.

https://securitystudio.com/top-7-vendor-related-breaches-of-all-time/

Evaluating vendors is vital for businesses to adhere to regulations, as these regulations often pertain to the vendors they collaborate with. Ensuring vendors comply with applicable regulations and industry standards protects the organization from fines and other legal consequences.

Additionally, vendor assessments provide evidence of due diligence and compliance checks, which are crucial during audits and investigations. Vendor assessment also promotes transparency and accountability. Organizations gain insight into vendor security capabilities by thoroughly evaluating vendor security practices. This knowledge supports better risk assessment and vendor selection. Furthermore, vendor assessments create a framework for monitoring and reviewing vendors' performance and security practices. Continuous evaluation helps ensure that vendors maintain their commitment to security and remain aligned with the organization.

Vendor Diversity

Vendor diversity is essential for several reasons, offering benefits not only in terms of cybersecurity but also in business resilience, innovation, and competition:

Cybersecurity — Relying on a single vendor for all software and hardware solutions can create a single point of failure. The entire infrastructure may be at risk if a vulnerability is discovered in that vendor's products. Vendor

diversity introduces multiple technologies, reducing the impact of a single vulnerability and making it more difficult for attackers to exploit the entire system.

Business resilience — Vendor diversity mitigates the risk associated with vendor lock-in and ensures that an organization's operations are not solely reliant on one vendor's products or services. If a vendor stops doing business, goes bankrupt, or experiences a significant disruption, having alternatives helps maintain business continuity.

Innovation — Diverse vendors bring different perspectives, ideas, and technologies. Leveraging solutions from multiple vendors can lead to a more innovative and agile IT infrastructure, better positioning an organization to adapt to emerging trends and technologies.

Competition — Vendor diversity promotes healthy competition in the market, which can lead to better pricing, improved product features, and higher-quality customer support. By engaging multiple vendors, organizations can encourage continuous improvement and obtain better value for their investments.

Customization and flexibility — Different vendors offer unique solutions that cater to specific needs, and having a diverse vendor ecosystem allows organizations to choose the best fit for their requirements. This flexibility can result in a more tailored and effective IT infrastructure.

Risk management — Vendor diversity helps spread the risk associated with potential product or service failures, security breaches, and other issues. Organizations can better manage and mitigate risks by not trusting a single solution provider or supplier.

Compliance — In some industries, regulations or industry standards may require organizations to maintain vendor diversity to ensure compliance and reduce the risk of supply chain disruptions or security breaches.

Conflict of Interest

A conflict of interest arises when an individual or organization has competing interests or obligations that could compromise their ability to act objectively, impartially, or in the best interest of another party. When performing vendor assessments, it is vital to determine whether a vendor's interests, relationships, or affiliations may influence their ability to provide unbiased recommendations, fair pricing, or deliver services without bias. Organizations must diligently identify and address potential conflicts of interest, including scrutinizing the vendor's affiliations, relationships with competitors or stakeholders, financial interests, and any potential bias that could compromise their integrity. Some examples of conflict of interest include the following items:

Financial interests — A vendor may have a financial interest in recommending specific products or services due to partnerships, commissions, or financial incentives that bias their recommendations and lead to selecting options that may not fit the organization's needs.

Personal relationships — If a vendor has personal relationships or close ties with decision-makers within the organization, it can influence decision-making and compromise the objective evaluation of other vendors.

Competitive relationships — A vendor may have a business relationship or competitive interest with another vendor under consideration, which can lead a vendor to prioritize their own interests or partnerships over the organization's best interests.

Insider information — In cases where a vendor has access to confidential or proprietary information about other vendors or the organization's strategic plans, the vendor may use this information to gain an unfair advantage or manipulate the selection process.

Vendor Assessment Methods

Due diligence, in the context of vendor assessment and selection, refers to the comprehensive and systematic process of gathering and analyzing information about potential vendors to assess their suitability, reliability, and integrity. It involves conducting a thorough investigation and evaluation of vendors based on predetermined criteria, including financial stability, reputation, technical capabilities, security practices, regulatory compliance, and past performance.

Due diligence aims to minimize risks and support informed decisions during the vendor selection process by verifying the accuracy of vendor claims, identifying potential red flags, and ensuring alignment with the organization's needs. Through due diligence, organizations can uncover any undisclosed risks or issues, clearly understand the vendor's capabilities and limitations, and evaluate the potential impact on business operations.

Assessment Method	Description
Penetration testing	Penetration testing evaluates vendors' security posture and identifies potential vulnerabilities in their systems, networks, and applications. By conducting penetration tests on vendor infrastructure or seeking evidence that penetration tests have been performed, organizations can gain insights into the vulnerabilities that attackers could exploit, helping them understand the potential risks associated with partnering with the vendor. Penetration testing provides a comprehensive assessment of the vendor's security resilience, allowing businesses to make informed decisions about their suitability as a vendor. Penetration tests improve the vendor assessment process by validating the effectiveness of security controls, uncovering hidden weaknesses, and assisting risk management practices.
Right-to-audit clause	A right-to-audit clause is a contractual provision that grants an organization the authority to conduct audits or assessments of vendor operational practices, information systems, and security controls. The right-to-audit clause supports vendor assessment practices by allowing organizations to validate and verify the vendor's compliance with contractual obligations, security standards, and regulatory requirements. By exercising the right to audit, organizations can gain transparency into the vendor's operations, identify gaps or deficiencies, and ensure that the vendor maintains the expected level of security and compliance at all times.
Evidence of internal audits	When performing vendor due diligence, looking for evidence that the vendor has internal audit practices is crucial. Internal audit provides an independent and objective evaluation of an organization's internal controls, risk management practices, and compliance with policies and regulations. By examining the presence and effectiveness of internal audits within a vendor's operations, businesses can gain confidence in the vendor's commitment to good governance, risk management, and compliance. Evidence of internal audit demonstrates that the vendor has established mechanisms for internal oversight, periodic assessments, and continuous improvement of their processes. It demonstrates a proactive approach to risk identification and mitigation and a commitment to secure operations.
Independent assessments	Organizations often rely on independent assessments as crucial vendor selection criteria. Independent assessments involve engaging with independent experts to evaluate and verify vendor capabilities, security, and compliance practices. These assessments provide an objective and unbiased evaluation of vendor capabilities. By leveraging independent assessments, organizations can benefit from specialized knowledge and industry best practice approaches to security assessments that internal teams may not know. Leveraging independent assessments helps mitigate potential biases, ensure thorough evaluations, and support informed decision-making during the vendor selection process. Additionally, periodic reassessment of existing vendors fosters continuous improvement in the vendor's security practices.
Supply chain analysis	Supply chain refers to the interconnected network of entities involved in producing, distributing, and delivering goods or services, from raw material suppliers to manufacturers, distributors, retailers,

Assessment Method	Description
	and ultimately, the end customer. It is a complex ecosystem involving collaboration between numerous vendors at various stages. Vendors are essential in the supply chain as they provide goods, services, and expertise contributing to the final product or service. Vendors often include raw materials suppliers, manufacturers, logistics providers, and technology solution providers. Each vendor within the supply chain has its own set of capabilities, processes, and potential risks. Managing and assessing these vendors is crucial to ensure the smooth flow of materials, minimize disruptions, maintain quality standards, and uphold security and compliance requirements throughout the supply chain. Supply chain analysis evaluates the risks and vulnerabilities associated with the various entities (vendors) involved in a supply chain by examining the security practices, capabilities, and reliability of individual vendors within the supply chain network and how security issues with one vendor in the chain may compromise the security of the organization's environment. This information helps organizations identify weak links, vulnerabilities, and potential points of compromise within the supply chain so they can be addressed before they cause problems.

Performing vendor site visits offers firsthand observation and assessment of a vendor's physical facilities, operational processes, and overall risk management practices, allowing for a more comprehensive evaluation of potential risks and vulnerabilities.

Vendor Monitoring

Vendor monitoring involves continuously overseeing and evaluating vendors to ensure ongoing adherence to security standards, compliance requirements, and contractual obligations. It may include regular performance reviews, periodic assessments, and real-time monitoring of vendor activities. This proactive approach allows organizations to identify and address potential risks or issues promptly.

Questionnaires

Questionnaires gather vendor information about their security practices, controls, and risk management strategies to help organizations assess a vendor's security posture, identify vulnerabilities, and evaluate their capabilities. Questionnaires provide a structured means of obtaining consistent vendor information, enabling more effective risk analysis and comparison fairly and consistently. Questionnaires collect information about the vendor's security policies, procedures, and controls, including data protection, access management, incident response, and disaster recovery. The questionnaire may ask about a vendor's compliance with industry-specific regulations and standards, such as GDPR, HIPAA, ISO 27001, or PCI-DSS. It may also seek details about the vendor's security training and awareness programs for employees and their approach to conducting third-party security assessments and audits. Additionally, the questionnaire may explore the vendor's incident response capabilities, breach history, and insurance coverage.

The answers to vendor risk management questionnaires should be validated by requesting supporting documentation, conducting site visits or audits, performing background checks, contacting references or previous clients, and utilizing third-party verification services to ensure the accuracy and reliability of the vendor's responses.

Rules of Engagement

Rules of Engagement (RoE) define the parameters and expectations for vendor relationships. These rules outline the responsibilities, communication methods, reporting mechanisms, security requirements, and compliance obligations that vendors must adhere to. Rules of engagement establish clear guidelines for the vendor's behavior, activities, and access to sensitive information. By setting these boundaries, organizations can establish a controlled and secure environment, mitigating the potential risks associated with third-party relationships. Some important elements included in an RoE include the following:

Roles and responsibilities — Clearly define the roles and responsibilities of the vendor and client in managing risks, including specifying who is responsible for identifying, assessing, and mitigating various types of risks.

Security requirements — Outline the security standards, practices, and controls the vendor must adhere to, including provisions related to data protection, access controls, encryption, incident response, and regular security assessments.

Compliance obligations — State the regulatory and compliance obligations the vendor must meet, ensuring they align with the client's industry-specific requirements, including privacy, data security, and any other applicable legal or industry regulations.

Reporting and communication — Establish protocols for timely reporting of security incidents, breaches, or potential risks, including defining the reporting channels, frequency, and level of detail required to ensure effective risk communication and management.

Change management — Outline procedures for managing changes or updates to systems, processes, or services that could impact security and introduce new risks, including change approval processes, testing requirements, and documentation practices.

Contractual provisions — Include provisions related to indemnification, liability, insurance, and termination rights in case of security breaches or failure to meet risk management obligations. These provisions help allocate responsibilities and provide legal recourse in case of noncompliance or breaches.

Legal Agreements

Legal agreements play a vital role in supporting vendor relationships by establishing both parties' rights, responsibilities, and expectations. Legal agreements serve as the foundation for the vendor-client relationship, providing a framework for conducting business and addressing potential issues or disputes that may arise.

Initial Agreements

Different types of agreements are needed to govern vendor relationships based on the specific nature of the engagement and the services being provided. The following agreements play distinct roles in setting up vendor relationships:

Memorandum of Understanding (MOU) — a nonbinding agreement that outlines the intentions, shared goals, and general terms of cooperation between parties. MOUs serve as a preliminary step to establish a common understanding before proceeding with a more formal agreement.

Non-disclosure Agreement (NDA) — ensures the confidentiality and protection of sensitive information shared during the relationship. An NDA is a binding agreement likely to be signed alongside an MOU.

Memorandum of Agreement (MOA) — a formal agreement that defines the parties' specific terms, conditions, and responsibilities. MOAs establish a legally binding relationship covering objectives, roles, resources, and obligations. They provide a trustworthy framework for collaboration.

Business Partnership Agreement (BPA) — governs long-term strategic partnerships between organizations. BPAs encompass various objectives, including goals, financial arrangements, decision-making processes, intellectual property rights, confidentiality, and dispute resolution mechanisms. BPAs provide a means for governing collaborative and mutually beneficial relationships.

Master Service Agreement (MSA) — outlines the overall terms and conditions of a specific contract, such as provisioning cloud resources or ticketing/help desk support. An MSA includes scope, pricing, deliverables, and intellectual property rights.

Detailed Agreements

Where initial agreements establish a framework for collaboration or service provision, other agreements can be implemented to specify terms for operational detail. These help to govern vendor relationships effectively.

Service Level Agreement (SLA) — defines the specific performance metrics, quality standards, and service levels expected from the vendor.

Statement of Work (SOW)/Work Order (WO) — details a vendor project or engagement's scope, deliverables, timelines, and responsibilities. SOWs clarify the vendor's tasks, the organization's expectations, and the agreed-upon deliverables. They are crucial for managing project execution and ensuring vendor and organization alignment.

12.2.3 Practice Questions (Section Quiz)

q_man_thirdparties_bpa_01_secp8

Two organizations plan on forming a partnership to provide systems security services. Onboarding requirements for both sides include a mutual understanding of quality management processes.

Which approach BEST meets this requirement?

Answers:

***Business partnership agreement (BPA)**

Non-disclosure agreement (NDA)

Service level agreement (SLA)

Measurement systems analysis (MSA)

Explanation:

BPAs are commonly used models in IT, such as among partner agreements that large IT companies set up with resellers and solution providers.

An NDA is an agreement that provides a basis for protecting information assets. NDAs are between companies and employees, between companies and contractors, and between two companies.

An SLA is a formal agreement that lays out the detailed conditions for how the vendor conducts the service.

An MSA relates to quality management processes that use quantified analysis methods to determine the effectiveness of a system and may be part of an onboarding requirement.

q_man_thirdparties_bpa_02_secp8

A new IT security firm is partnering with an IT support company and is opening its business soon. The firm would like to be a reseller for a popular firewall.

Which of the following options allows the firm to become an authorized reseller?

Answers:

***Business partners agreement (BPA)**

Memorandum of agreement (MOA)

Memorandum of understanding (MOU)

Non-disclosure agreement (NDA)

Explanation:

A BPA is a partner agreement that large IT companies set up with resellers and solution providers.

An MOA is a formal agreement, or contract, that contains specific obligations rather than a broad understanding.

An MOU is a preliminary or exploratory agreement to express an intent to work together. MOUs tend to be relatively informal and do not act as binding contracts.

An NDA is an agreement that provides a basis for protecting information assets. NDAs are agreements between companies and employees, between companies and contractors, or between two companies.

q_man_thirdparties_discuss_differences_secp8

As a new IT manager at TechCorp, you are tasked with onboarding a third-party vendor that will provide critical IT services.

During the onboarding process, you discover that the vendor's security policies and incident response procedures are significantly different from those of TechCorp.

What should you do?

Answers:

Ignore the differences and proceed with the onboarding process.

Cancel the onboarding process immediately.

***Discuss the differences with the vendor and seek to align the policies and procedures.**

Report the vendor to the authorities for having different policies.

Explanation:

When differences in security policies and procedures are identified, the best course of action is to discuss these differences with the vendor. The goal should be to align the policies and procedures as closely as possible to ensure the security of both organizations.

Ignoring significant differences in security policies and procedures could expose TechCorp to unnecessary risks. It's important to address these differences before proceeding with the onboarding process.

While it's important to take security seriously, cancelling the onboarding process immediately doesn't allow for the possibility of resolving the differences. It's better to discuss the issues with the vendor first.

Having different security policies and procedures is not illegal, and there's no need to report the vendor to the authorities. The focus should be on working with the vendor to align the policies and procedures.

q_man_thirdparties_insider_information_secp8

You are the procurement manager in a large corporation. Your company is in the process of selecting a new software vendor.

One of the vendors in the running is a company where your cousin is a senior executive. Your cousin has shared some insider information about their upcoming product updates that could give them an edge over the competition.

What should you do?

Answers:

Use the insider information to influence the vendor selection process in favor of your cousin's company.

***Disclose the relationship and the information received to your supervisor and remove yourself from the decision-making process.**

Keep the information to yourself and continue with the vendor selection process.

Share the insider information with the other vendors to level the playing field.

Explanation:

Disclosing the relationship and the information received is the best choice as it ensures transparency and removes potential bias from the decision-making process. It is the most ethical course of action and aligns with best practices for managing conflicts of interest.

Using insider information to influence a business decision is a conflict of interest and could lead to legal consequences.

Keeping the information to yourself and continuing with the vendor selection process could lead to biased decisions and potential legal consequences.

Sharing insider information with other vendors is a breach of confidentiality and could lead to legal consequences. It also does not address the conflict of interest.

q_man_thirdparties_mou_01_secp8

You are the IT manager at TechCorp and are in the process of establishing a new partnership with a third-party vendor, VendorX, for a critical IT project.

The project involves data sharing and system integration between TechCorp and VendorX. As part of the onboarding process, you need to create an interoperability agreement (IA).

Which of the following documents should be included in the IA to specify who is going to do what, and when they will do it?

Answers:

Service level agreement (SLA)

Non-disclosure agreement (NDA)

***Memorandum of understanding (MOU)**

Blanket purchase order (BPO)

Explanation:

A memorandum of understanding (MOU) is the correct answer. An MOU provides a summary of which party in the relationship is responsible for performing specific tasks. In other words, the MOU specifies who is going to do what, and when they will do it.

A service level agreement (SLA) is important for defining the services performed by the third party and the level of performance guaranteed. However, it does not typically specify who is going to do what and when they will do it.

A non-disclosure agreement (NDA) is a contract in which the third party agrees not to share any of the information gathered during the completion of the work. While it's an important part of the IA, it does not specify who is going to do what and when they will do it.

A blanket purchase order (BPO) is an agreement with a third-party vendor to provide services on an ongoing basis. While it's a critical part of the IA, it does not specify who is going to do what and when they will do it.

q_man_thirdparties_mou_02_secp8

Two technology firms are in preliminary discussions to work together on several projects. The joint venture's goal entails providing support services to a broader customer base as an entity with shared resources.

Each firm has its own customer base, custom-branded products, and established processes.

Which of the following types of agreements BEST meets the firms' needs?

Answers:

***Memorandum of understanding (MOU)**

Memorandum of agreement (MOA)

Business partners agreement (BPA)

Non-disclosure agreement (NDA)

Explanation:

An MOU is a preliminary or exploratory agreement to express an intent to work together. MOUs tend to be relatively informal and do not act as binding contracts.

An MOA is a formal agreement or contract that contains specific obligations rather than a broad understanding.

A BPA is a type of partner agreement that large IT companies, such as Microsoft and Cisco, set up with resellers and solution providers.

An NDA is an agreement that provides a basis for protecting information assets. NDAs can exist between companies and employees, between companies and contractors, and between two companies.

q_man_thirdparties_nda_01_secp8

The IT department of a medium-sized company is in the process of finalizing agreements with various vendors. The legal team drafted the contracts to ensure proper arrangements.

The team considers three types of agreements: an NDA, a BPA, and an MOU. The IT team wants to select the MOST appropriate agreement for each vendor to ensure smooth collaboration.

Which of the following agreements protects sensitive information shared between the company and its vendors?

Answers:

***Non-disclosure agreement (NDA)**

Business partnership agreement (BPA)

Memorandum of understanding (MOU)

Memorandum of agreement (MOA)

Explanation:

The non-disclosure agreement (NDA) is a suitable agreement for protecting sensitive information shared between parties and maintaining confidentiality.

The business partnership agreement (BPA) centers on defining the terms and conditions of a partnership without a specific focus on protecting sensitive information.

Although a memorandum of understanding (MOU) outlines the intentions and expectations of parties, it does not deal specifically with information confidentiality.

A memorandum of agreement (MOA) is a formal agreement or contract that contains specific obligations rather than a broad understanding and does not emphasize maintaining confidentiality between parties.

q_man_thirdparties_nda_02_secp8

In a technology company, the IT department seeks to establish a new partnership with a cloud service provider. The IT team wants well-defined partnership terms and adequate protection of sensitive data shared between the company and the vendor.

The team is considering various agreement types to achieve this goal.

What type of agreement should the IT department consider to protect sensitive information shared between the technology company and the cloud service provider?

Answers:

***Non-disclosure agreement (NDA)**

Memorandum of understanding (MOU)

Business partnership agreement (BPA)

Service level agreement (SLA)

Explanation:

Parties ensure the confidentiality and protection of sensitive information shared through the non-disclosure agreement (NDA). NDAs establish clear guidelines on what information both parties consider confidential and outline their obligations and responsibilities in safeguarding this data.

Parties actively outline their intentions and expectations in a potential partnership or collaboration through a memorandum of understanding (MOU).

Partners actively cover various aspects of the partnership, such as roles, responsibilities, decision-making processes, and dispute resolution in a business partnership agreement (BPA).

Parties enter an active contractual agreement known as the service level agreement (SLA), which details factors like performance metrics, response times, and penalties for non-compliance.

q_man_thirdparties_nda_03_secp8

The IT department in a technology company is finalizing an agreement with a cloud service provider to host sensitive customer data. The company's legal team is drafting the contract, which includes a service level agreement (SLA) and a non-disclosure agreement (NDA).

Which of the following explanations MOST accurately demonstrates the primary purpose of including an NDA in the contract with the cloud service provider?

Answers:

***To protect the confidentiality of the company's data and proprietary information**

To specify the expected service quality and support responsiveness

To ensure compliance with industry regulations and standards

To outline the vendor's responsibilities for incident response and recovery

Explanation:

Integrating an NDA into the contract protects the company's sensitive data and unique proprietary knowledge. This agreement forms a legal foundation that keeps this information secure and prevents unauthorized entities from inadvertently or maliciously disclosing it.

In contrast to the NDA, the SLA sets out the expected level of service that we expect the vendor to deliver, including standards for uptime and the speed of support responses. This ensures that the vendor meets our high service standards.

The NDA maintains the data's confidentiality and protects it from potential breaches.

The NDA strengthens our focus on confidentiality and establishes strong data protection measures.

q_man_thirdparties_nda_04_secp8

A popular entertainment company is onboarding a new employee. The company has completed preliminary interview steps and due diligence.

Internal security is extremely important, so their human resources department is preparing documentation for the formal employment process.

In implementing the process, which solution would help limit the risk of proprietary data that an employee outside the company can use?

Answers:

***Non-disclosure agreement (NDA)**

Identity and access management (IAM)

Background check

Analysis and identification

Explanation:

When an employee or contractor signs an NDA, they confirm they will not share confidential information with a third party. Signing this type of contract legally protects internal intellectual property.

IAM refers to creating an account for the user to access the computer system with the appropriate privileges.

The background check process essentially determines that a person is who they say they are and are not concealing criminal activity, bankruptcy, or connections that would make them unsuitable or risky employees.

Analysis and identification are a part of the incident lifecycle. During this phase, the identification of a threat can occur.

q_man_thirdparties_right-to-audit_secp8

You are the chief information security officer (CISO) at a large financial institution. Your company is considering a new third-party vendor for a critical data processing service.

The vendor has a strong reputation and offers a cost-effective solution. However, during the assessment, you find that the vendor's security practices are not fully transparent.

What should you do?

Answers:

Proceed with the vendor since they are cost-effective and have a strong industry reputation.

Reject the vendor immediately due to their lack of transparency.

***Negotiate a right-to-audit clause in the contract to ensure you can assess their security practices on an ongoing basis.**

Ignore the lack of transparency since the vendor's solution is critical to your operations.

Explanation:

Negotiating a right-to-audit clause in the contract is the best choice as it allows you to assess the vendor's security practices on an ongoing basis. This ensures that the vendor maintains appropriate security controls and adheres to your company's security standards.

While cost-effectiveness and reputation are important, they should not override security concerns, especially when dealing with critical data processing services.

While rejecting the vendor immediately due to their lack of transparency might seem like a safe choice, it's not necessarily the best initial action. Rejecting the vendor immediately does not leave room for potential improvements or negotiations that could lead to a win-win situation.

Ignoring the lack of transparency could lead to significant vulnerabilities and potential breaches. This could have severe consequences for your company.

q_man_thirdparties_rules_of_engagement_secp8

You are a cybersecurity consultant hired to conduct a penetration test for a client. The client has provided you with a rules of engagement (RoE) document.

Upon reviewing the document, you notice that it does not specify the timeframes for the testing activities.

What should you do?

Answers:

Proceed with the penetration test at your convenience since the RoE does not specify timeframes.

Reject the RoE and refuse to conduct the penetration test until timeframes are specified.

***Engage the client in a discussion to clarify and agree upon the timeframes for the testing activities.**

Decide on the timeframes yourself and inform the client after you have started the penetration test.

Explanation:

Engaging the client in a discussion to clarify and agree upon the timeframes is the best choice as it ensures that both parties are on the same page. This is an essential part of the RoE and should be clearly defined before the penetration test begins.

Proceeding with the penetration test without agreed-upon timeframes could lead to disruptions in the client's operations and potential misunderstandings.

Rejecting the RoE and refusing to conduct the penetration test until timeframes are specified might seem like a safe choice, but it's not necessarily the best initial action. Rejecting the RoE immediately does not leave room for potential improvements or negotiations that could lead to a win-win situation.

Deciding on the timeframes yourself and informing the client after you have started the penetration test could lead to disruptions and potential misunderstandings. It's important to agree on the timeframes before the penetration test begins.

q_man_thirdparties_sla_01_secp8

Which of the following is defined as a contract that prescribes the technical support or business parameters a provider bestows to its client?

Answers:

Mutual aid agreement

***Service level agreement**

Final audit report

Certificate practice statement

Explanation:

A service level agreement is defined as a contract that prescribes the technical support or business parameters a provider bestows to its client.

A mutual aid agreement is an agreement between two organizations to support each other in the event of a disaster.

A final audit report is the result of an external auditor's inspection and analysis of an organization's security status.

A certificate practice statement defines the actions and promises of a certificate service authority.

q_man_thirdparties_sla_02_secp8

What is a service level agreement (SLA)?

Answers:

An agreement to support another company in the event of a disaster.

A contract with a legal entity to limit your asset-loss liability.

***A guarantee of a specific level of service.**

A contract with an ISP for a specific level of bandwidth.

Explanation:

An SLA is a guarantee of a specific level of service from a vendor. That service may be communication links, hardware, or operational services. An SLA is a form of insurance against disasters or security intrusions that may affect your organization's mission-critical business functions.

An agreement to support another company in the event of a disaster is known as a mutual aid agreement.

A contract with a legal entity to limit your asset-loss liability is an insurance policy.

A contract with an ISP for a specific level of bandwidth is a service contract.

q_man_thirdparties_sla_03_secp8

An organization is considering a hybrid cloud deployment to leverage the benefits of both private and public cloud resources.

While reviewing third-party vendors, what critical aspect should the employees consider for a secure and effective transition?

Answers:

***Establish clear service level agreements (SLAs)**

Prioritize lowest-cost vendors

Focus on data redundancy

Delegate all security management to the provider

Explanation:

Clear SLAs are vital for defining performance, availability, and support expectations. As the organization chooses the correct vendor, it must oversee the establishment and execution of service level agreements as a quality control measure.

While cost is an important factor in decision making, it should not take precedence over security considerations. Choosing a vendor based solely on cost can compromise security and compliance standards.

Although data redundancy is important in maintaining data availability and preventing data loss, it is not the only security consideration.

While third-party vendors can handle parts of the security management, the organization needs to stay involved.

q_man_thirdparties_sla_04_secp8

When performing forensic investigation in public clouds, what document would contain the right-to-audit clause and give the investigator the authority to audit files on the network?

Answers:

***Service level agreement (SLA)**

Forensic reports

Supply chain analysis

Checksums

Explanation:

An SLA is a formal agreement that lays out the detailed conditions for how the vendor will conduct the service. These could include terms and conditions for security access controls and risk evaluations, plus authentication criteria for proprietary and private data.

A digital forensics report summarizes the significant contents of digital data and the conclusions of the analysis.

Supply chain analysis evaluates the risks and vulnerabilities associated with the various entities (vendors) involved in a supply chain by examining the security practices, capabilities, and individual vendors' reliability within the supply chain network.

A checksum of two hashes can validate that the investigators did not tamper with the contents of a disk.

q_man_thirdparties_sla_05_secp8

At a technology company, the IT department is finalizing an agreement with a cloud service provider to host its sensitive customer data. The IT team has actively ensured the inclusion of a service level agreement (SLA) in the contract.

What is the primary purpose of actively including an SLA in the contract with the cloud service provider?

Answers:

***To define the level of service the cloud service provider must deliver.**

To protect the confidentiality of sensitive information shared between parties.

To outline the intentions and expectations of parties involved in a potential partnership.

To establish clear guidelines on what information is considered confidential.

Explanation:

The cloud service provider actively includes an SLA to define the level of service it must deliver, covering performance metrics, response times, and availability.

While a non-disclosure agreement (NDA) holds importance in protecting sensitive information, it is not the primary purpose of the SLA inclusion in the contract.

An SLA's primary purpose is not to outline the intentions and expectations of parties in a partnership. This is the role of a memorandum of understanding (MOU).

Establishing guidelines on confidential information is not the primary purpose of an SLA; this responsibility falls under the purview of an NDA.

q_man_thirdparties_supply_chain_analysis_secp8

You are the chief procurement officer in a multinational corporation. Your company is considering a new vendor for a critical component of your product. The vendor has a strong reputation and their product is of high quality.

However, you are aware that the vendor relies heavily on a single supplier for their raw materials.

What should you do?

Answers:

Proceed with the vendor since they have a strong reputation and their product is of high quality.

Reject the vendor immediately due to their dependency on a single supplier.

***Conduct a thorough supply chain analysis to assess the potential risks associated with the vendor's dependency on a single supplier.**

Negotiate lower prices with the vendor due to their dependency on a single supplier.

Explanation:

Conducting a thorough supply chain analysis is the best choice as it allows you to assess the potential risks associated with the vendor's dependency on a single supplier. This could include assessing the supplier's reliability, financial stability, and contingency plans. This information will help you make an informed decision.

While the vendor's reputation and product quality are important, they should not override potential supply chain risks, especially when dealing with critical components.

While rejecting the vendor immediately due to their dependency on a single supplier might seem like a safe choice, it's not necessarily the best initial action. Rejecting the vendor immediately does not leave room for potential improvements or negotiations that could lead to a win-win situation.

Negotiating lower prices does not address the potential supply chain risks. If the supplier fails to deliver, it could disrupt the vendor's operations and, in turn, your own.

q_man_thirdparties_vendor_assessment_secp8

As a cybersecurity manager, you are tasked with assessing a new third-party vendor that your company is considering for a critical software solution.

The vendor has a strong reputation in the industry and offers a cost-effective solution. However, during the assessment, you find that the vendor's security practices do not fully align with your company's security standards.

What should you do?

Answers:

Proceed with the vendor since they are cost-effective and have a strong industry reputation.

Reject the vendor immediately and look for other options.

***Engage the vendor in a discussion about your security concerns and see if they are willing to improve their practices.**

Ignore the security concerns since the vendor's solution is critical to your operations.

Explanation:

Engaging the vendor in a discussion about the security concerns is the best choice as it shows due diligence and opens the door for potential improvements. If the vendor is willing to align their practices with your security standards, it could lead to a beneficial partnership.

While cost-effectiveness and reputation are important, they should not override security concerns, especially when dealing with critical software solutions. This could expose the company to significant risks.

Rejecting the vendor immediately and looking for other options might seem like a safe choice, but it's not necessarily the best initial action. Rejecting the vendor immediately does not leave room for potential improvements or negotiations that could lead to a win-win situation.

Ignoring security concerns, especially for a critical software solution, could lead to significant vulnerabilities and potential breaches. This could have severe consequences for the company.

q_man_thirdparties_vendor_diversity_secp8

You are the chief information officer (CIO) of a large corporation. Your company has been relying on a single vendor for its entire IT infrastructure for the past five years.

Recently, this vendor has been facing financial difficulties and there are rumors of a potential bankruptcy.

What should be your immediate course of action?

Answers:

Ignore the rumors and continue business as usual.

Immediately switch to another vendor to avoid potential disruptions.

***Start diversifying your vendor portfolio to ensure business resilience.**

Negotiate lower prices with the current vendor due to their financial difficulties.

Explanation:

Diversifying your vendor portfolio ensures business resilience is the best choice. It mitigates the risk associated with vendor lock-in and ensures that your company's operations are not solely reliant on one vendor's products or services. If a vendor stops doing business, goes bankrupt, or experiences a significant disruption, having alternatives helps maintain business continuity.

Ignoring potential risks, especially when they could affect the entire IT infrastructure of your company, is not a good practice. This could lead to significant disruptions in case the vendor goes bankrupt.

Immediately switching to another vendor to avoid potential disruptions might seem like a safe choice, but it could lead to immediate disruptions and potential compatibility issues. Switching to a new vendor without proper planning and assessment could introduce new risks.

While negotiating lower prices might seem like a good short-term solution, it does not address the potential risk of the vendor going bankrupt. This could still lead to significant disruptions in your IT infrastructure.

12.3 Audits and Assessments

As you study this section, answer the following questions:

What is an audit?

What are the different types of audits?

How do the types of audits differ from one another?

In this section, you will learn to:

Audit the Windows security log.

Configure advanced audit policies.

Audit device logs on a switch.

Enable device logs.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	4.3 Explain various activities associated with vulnerability management. Identification methods System/process audit Validation of remediation Audit 4.4 Explain security alerting and monitoring concepts and tools. Activities Log aggregation 4.5 Given a scenario, modify enterprise capabilities to enhance security. Operating system security Group Policy 4.9 Given a scenario, use data sources to support an investigation. Log data OS-specific security logs Network logs 5.1 Summarize elements of effective security governance. Policies Information security policies 5.5 Explain types and purposes of audits and assessments. Attestation Internal Compliance

	Audit committee Self-assessments External Regulatory Examinations Assessment Independent third-party audit
TestOut Security Pro	1.2 Harden Authentication 1.2.5 Configure and Link Group Policy Objects (GPO) 5.1 Implement logging and auditing 5.1.1 Configure advanced audit policy

12.3.1 Audits (Lesson Video)

Transcript:

Audits and assessments are crucial to maintaining trustworthy relationships and business operations. Audits involve evaluating an organization's processes, controls, and compliance, whereas an assessment evaluates operations such as cybersecurity, risk management, and internal controls. In this lesson, we'll look at the difference between internal and external assessments. We'll also look at penetration testing, one of the more common assessments that need to be performed.

When an audit or assessment is done, it'll be classified as either an internal or external assessment. An internal audit is performed by the organization's employees to provide an in-depth analysis of the organization's business processes. These internal audits should be performed on a regular basis and focus on the needs and priorities of the organization. External audits, on the other hand, are performed by an independent third party. An external audit ensures that the organization's practices meet industry standards and identifies any areas the internal audit might have missed. Performing internal and external audits and assessments provides a full accounting of the organization's policies and procedures and allows the organization to quickly identify and fix any shortcomings. Performing these audits and assessments fosters transparency and accountability and builds trust among stakeholders. One of the more common assessments that should be performed on a regular basis is a penetration test.

Penetration testing, also called pen testing or ethical hacking, involves hacking into a network to discover any vulnerabilities and weaknesses so they can be remediated before they're exploited by a malicious hacker. The penetration test utilizes the same techniques that a hacker would use, but the pen tester has written permission to perform the audit, and the scope of the test is clearly defined before anything is done.

A penetration test will typically consist of five steps. These include reconnaissance, scanning, vulnerability assessment, exploitation, and reporting. In the reconnaissance phase, the pen tester begins gathering information about their target. This can include gathering publicly available information, using social engineering techniques, and dumpster diving.

The second step is scanning the target. At this point, the tester is actively engaged with the target. Various tools are used to gather in-depth information about the network, such as live hosts, open ports, and more.

During the vulnerability assessment phase, the pen tester takes all the information they've gathered so far and identifies all potential vulnerabilities that may exist on the network. The pen tester then moves into the next phase and attempts to exploit the confirmed vulnerabilities.

The final step, reporting, involves compiling a detailed analysis of the organization's security systems. This includes all discovered vulnerabilities that were exploited, along with recommendations to remediate the vulnerabilities.
When defining the rules of the penetration test, the first step is to determine whether an unknown, partially known, or known test environment will be used. Let's look at the difference in each of these.
An unknown test is when an ethical hacker has no information about the target or network. This test is the best for simulating an outside attack and ignores insider threats. The major drawback to this type of test is that it takes more time and is much more expensive, mostly because a lot of time and resources must go into the reconnaissance and scanning phases.
A known test is the opposite of an unknown test. The pen tester is given full knowledge of the network, computer systems, and infrastructure. The known test environment allows a comprehensive and thorough test of the organization's security. Still, it's not very realistic, as an attacker will rarely have all the information handed over to them.
The partially known testing environment simulates an insider threat. The tester is given partial information about the network and computer systems. This can be IP configurations, email lists, computer names, or other information an insider would realistically have. This type of test requires less reconnaissance and scanning but doesn't give all the information to the pen tester.
Aside from testing network security, physical penetration testing should also be performed. In this type of test, the penetration tester attempts to gain access to restricted areas, sensitive information, or critical assets within the organization using techniques like social engineering, tailgating, lock picking, bypassing alarms or surveillance systems, and exploiting physical vulnerabilities.
Red team and blue team are other terms you'll run across. In the cybersecurity field, an offensive security team is called a red team, and a defensive security team is called a blue team. As the red team attempts to break into a system, the blue team works to keep them out.
That'll wrap up this lesson on audits and assessments. In this lesson, we first looked at the benefits of performing audits and assessments and the difference between internal and external audits. We then went over penetration testing, including the steps involved and the different types of penetration tests that can be performed.

12.3.2 Audit Facts

Audits and assessments are crucial to maintaining trustworthy operations. Audits involve systematically evaluating processes, controls, and compliance with established standards, policies, and regulations. They ensure that an organization's operations align with defined requirements, identify gaps, and provide recommendations for improvement.

On the other hand, assessments involve evaluating the effectiveness and efficiency of various aspects of an organization's operations, such as cybersecurity, risk management, and internal controls. They help identify vulnerabilities, assess risks, and provide insights for enhancing security measures. Both audits and assessments play a vital role in maintaining compliance, mitigating risks, and continuously improving an organization's overall security and operational performance.

This lesson covers the following topics:

- Attestation and assessments
- Internal and external assessments
- Internal assessments
- External assessments

Attestation and Assessments

Attestation refers to verifying and validating the accuracy, reliability, and effectiveness of security controls, systems, and processes implemented within an organization. It involves an independent and objective examination by a qualified and trusted entity, such as an auditor or assessor. Attestation is a formal declaration or confirmation that an organization's security controls and practices comply with specific standards, regulations, or best practices and provides assurance to stakeholders, such as

management, customers, business partners, and regulators, that an organization's security measures are adequate and effective in protecting sensitive information, mitigating risks, and maintaining data confidentiality, integrity, and availability.

Internal and External Assessments

Using internal and external audit and assessment methods is essential for a comprehensive and effective evaluation of an organization's systems, controls, and management processes. The organization's employees conduct internal audits and provide an in-depth assessment of the organization's business processes. Internal teams can conduct regular, focused assessments that align with the organization's needs and priorities and support continuous monitoring and improvement of internal controls, governance and risk management practices, and operational efficiency.

In contrast, independent third-party service providers conduct external audits and assessments that utilize specialized expertise and knowledge in specific domains, regulations, and industry best practices. External auditors convey an impartial and objective evaluation of business practices that is impossible to obtain using internal teams. External audits ensure that the organization's practices are measured against recognized industry standards and help identify improvement areas that internal audit teams may have missed.

Organizations can achieve several important objectives by utilizing internal and external audit and assessment methods. Using both approaches helps facilitate a balanced and comprehensive view of the organization's risk management practices, controls, and compliance efforts. Combining internal and external audits enhances the organization's risk management capabilities. Internal audits enable continuous monitoring, early detection of issues, and timely remediation, while external audits validate the organization's controls, compliance, and risk mitigation efforts.

Additionally, utilizing both internal and external methods fosters transparency and accountability. Internal audits promote a culture of self-assessment and continuous improvement within the organization, while external audits provide stakeholders with independent assurance and validation of the organization's practices. This combination helps build trust among stakeholders, including customers, business partners, regulatory bodies, and investors. An often overlooked benefit, the collaboration between internal and external auditors facilitates knowledge sharing and professional development, improving the quality of both teams' assessments. Internal auditors can learn from the expertise and best practices of external auditors, while external auditors gain a deeper understanding of the organization's operating environment and the challenges they face that often impede compliance initiatives.

Internal Assessments

Internal assessments are required for government agencies according to the NIST RMF, PCI-DSS, and others.

Approach	Description
Compliance assessment	Internal compliance assessments ensure operating practices align with laws, regulations, standards, policies, and ethical requirements. These assessments evaluate the effectiveness of internal controls, identify noncompliance or risk areas, and communicate findings to stakeholders such as risk managers.
Audit committee	Audit committees provide independent oversight and assurance regarding an organization's financial reporting, internal controls, and risk management practices. These committees are typically composed of board members independent of the organization's management team. Audit committees aim to enhance the integrity of financial statements, ensure compliance with legal and regulatory requirements, monitor the effectiveness of internal controls, oversee the external audit process, and promote transparency and accountability. Audit committees are critical in fostering confidence among shareholders, stakeholders, and the public by providing an independent and objective assessment of the organization's financial practices and contributing to sound corporate governance.

Approach	Description
Self-assessment	Self-assessments allow individuals or organizations to evaluate their performance, practices, and adherence to established criteria against predetermined metrics and measures. Self-assessments help identify strengths, weaknesses, and areas for improvement, enabling individuals or organizations to take proactive measures to enhance their effectiveness and outcomes. Self-assessments imply internal personnel with the expertise, knowledge, and understanding of the assessed area are available to complete them.

Internal assessments are required for government agencies according to the NIST RMF, PCI-DSS, and others.

External Assessments

External entities could include certified public accountants (CPAs), external auditors, consulting firms, regulatory bodies, or specialized assessment agencies. The independence of these external assessors ensures impartiality and objectivity in the evaluation process.

Approach	Description
Regulatory	Regulatory authorities or agencies perform assessments to ensure compliance with specific laws, regulations, or industry standards. Regulatory assessments evaluate whether organizations adhere to mandatory regulatory requirements and promote a culture of compliance. Regulatory assessments typically involve inspections, audits, or reviews of processes, practices, and controls to verify compliance, identify deficiencies, and enforce regulatory obligations. Regulatory assessments play a critical role in safeguarding public interests, protecting consumers, maintaining market integrity, and upholding industry standards. They help mitigate risks, ensure fair competition, and enhance transparency and accountability in regulated industries.
Examination	An external examination typically refers to an independent and formal evaluation conducted by external parties, such as auditors or regulators, to assess the accuracy, reliability, and compliance of an organization's financial statements, processes, controls, or specific aspects of its operations. External examinations focus on verifying information accuracy and ensuring compliance with applicable laws, regulations, or industry standards. Examples of external examinations include financial statement audits, regulatory compliance audits, and specific assessments of control environments.
Assessment	An external assessment generally refers to a broad evaluation conducted by external experts or consultants to assess an organization's overall performance, practices, capabilities, or specific focus areas. External assessments can encompass various elements, such as strategy, operational efficiency, risk management, cybersecurity, or compliance practices. The goal is to provide an objective and independent perspective on the organization's strengths, weaknesses, and opportunities for improvement.
Independent third-party audit	Independent third-party audits provide objective and unbiased assessments of an organization's systems, controls, processes, and compliance. The importance of independent third-party audits lies in their ability to offer an external perspective, free from any conflicts of interest or bias. Independent audits instill confidence among stakeholders, including customers, business partners, regulatory bodies, and investors, as they attest to an organization's commitment to quality, compliance, and good governance. They also help organizations demonstrate transparency, accountability, and adherence to industry standards and regulations.

External entities could include certified public accountants (CPAs), external auditors, consulting firms, regulatory bodies, or specialized assessment agencies. The independence of these external assessors ensures impartiality and objectivity in the evaluation process.

12.3.3 Auditing the Windows Security Log (Demo Video)

Transcript:

In this demonstration, we'll show you how to audit the Windows Security log, allowing you to audit security-related events on one host or across your entire Windows network.

There are two methods for auditing:
You can audit individual workstations using the local security policy. However, this method can be time-intensive and challenging to manage when auditing more than a few hosts.
Alternatively, you can use Group Policy, which we'll opt for as it's associated with the entire Windows domain. This is the method we'll demonstrate, as it's a more efficient solution. When a domain user logs in at a workstation, the audit settings for the Windows Security log will be applied to that workstation.
Let's open Group Policy Management.
This is our default domain, and here we have the Default Domain policy.
We'll right-click and go to Edit. Now, we have a Group Policy Management Editor, which we'll expand.
We'll navigate to Computer Configuration, Policies, Windows Settings, Security Settings, Local Policies, and Audit Policy.
On the right pane, you'll see audit policies, a total of nine.
For instance, if we explore Audit Account Management, we can specify auditing the management of user and group accounts, or we can audit directory service access and various other events.
Object Access allows us to audit access to files, folders, printers, registry keys, and more.
We also have Audit Policy Change, which audits changes made to local policies, including user rights or security options.
Audit Privilege Use lets us audit the use of user rights or privileges.
Audit Process Tracking, primarily used by application developers, tracks detailed program information for events like program activation, process, exit, handle duplication, etc. It's resource-intensive, so use it sparingly.
Lastly, we have Audit System Events, which logs system events such as server reboots, providing valuable information.
Among the more useful audits from a security perspective are Audit Account Logon and Audit Logon Events. By default, these policies are undefined. Let's configure some Audit Logon Event policy settings.
Let's open this up. We can audit for both success and failure, depending on your preferences and security needs.
For example, auditing failed logon attempts can help identify potential security threats like password guessing attacks.
On the other hand, auditing successful logins can be useful when investigating unauthorized access.
For our demonstration, we'll audit logon failure events by selecting Failure and deselecting Success.
Now, let's close the Group Policy Management Editor. The Default Domain policy is linked to the domain. To apply these settings, we need to generate some failed logon attempts and analyze them using Event Viewer.
We've restarted this machine to update the Group Policy settings.
Instead of rebooting, you can also use the Search box and type "GPupdate" to update Group Policy.
Before logging back in, we intentionally created some failed password attempts. Let's examine the results.
In Event Viewer, right-click and run as an administrator, expand Windows logs, and go to Security.
We can see that our policy in the Default Domain has generated several audit failures due to the failed logon attempts.
We can view details such as the account name that attempted to log in with a bad password.
By navigating through the logs, we can inspect these failed attempts. Now, let's close this window.
Before concluding this demonstration, let's explore one more feature.
Under Group Policy Management, we'll return to the Default Domain, Edit, Group Policy Management Editor, and go to Computer Configuration, Windows settings.
We'll expand and go to Security Settings, Local Policies, and Audit Policies.
While the basic audit policies offer settings for broader behaviors, advanced audit policies allow for more granular control over auditing specific security events.
For instance, we can configure audit policies for logon or logoff events. Under Security Settings, Advanced Audit Policy Configuration, and Audit Policies, we can target more specific audit events.
Let's take Account Logon as an example. We can audit both success and failure of logon events.

With the earlier audit settings, a single logon failure generated multiple messages in the Windows Security Log. By selecting specific audit options here, we can significantly reduce the volume of information generated, making it more manageable.
These advanced audit policies also allow auditing for specific events like account lockouts, which isn't possible with the earlier policies.
In summary, that's it for this demo. We discussed setting up auditing, configured logon auditing, generated and examined failed logon events in Event Viewer, and delved into the Advanced Audit Policy Configuration for more granular security event auditing.

12.3.4 Configure Advanced Audit Policy (Simulation)

Scenario

You work as the IT security administrator for a small corporate network. As part of an ongoing program to improve security, you want to implement an audit policy for all workstations. You plan to audit user logon attempts and other critical events.

In this lab, your task is to configure the following audit policy settings in WorkstationGPO:

Local Policies	Setting
Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings	**Enabled**
Audit: Shut down system immediately if unable to log security audits	**Enabled**

Event Log	Setting
Retention method for security log	***Define: Do not overwrite events (clear log manually)***

Advanced Audit Policy Configuration	Setting
Account Logon: Audit Credential Validation	**Success and Failure**
Account Management: Audit User Account Management	**Success and Failure**
Account Management: Audit Security Group Management	**Success and Failure**
Account Management: Audit Other Account Management Events	**Success and Failure**
Account Management: Audit Computer Account Management	**Success**
Detailed Tracking: Audit Process Creation	**Success**
Logon/Logoff: Audit Logon	**Success and Failure**
Logon/Logoff: Audit Logoff	**Success**

Policy Change: Audit Authentication Policy Change	**Success**
Policy Change: Audit Audit Policy Change	**Success and Failure**
Privilege Use: Audit Sensitive Privilege Use	**Success and Failure**
System: Audit System Integrity	**Success and Failure**
System: Audit Security System Extension	**Success and Failure**
System: Audit Security State Change	**Success and Failure**
System: Audit IPsec Driver	**Success and Failure**

Do not use the old audit policies located in **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** > **Audit Policies** .

Explanation

While completing this lab, use the following WorkstationGPO settings:

Local Policies	**Setting**
Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings	**Enabled**
Audit: Shut down system immediately if unable to log security audits	**Enabled**

Event Log	**Setting**
Retention method for security log	***Define: Do not overwrite events (clear log manually)***

Advanced Audit Policy Configuration	**Setting**
Account Logon: Audit Credential Validation	**Success and Failure**
Account Management: Audit User Account Management	**Success and Failure**
Account Management: Audit Security Group Management	**Success and Failure**
Account Management: Audit Other Account Management Events	**Success and Failure**
Account Management: Audit Computer Account Management	**Success**
Detailed Tracking: Audit Process Creation	**Success**

Logon/Logoff: Audit Logon	**Success and Failure**
Logon/Logoff: Audit Logoff	**Success**
Policy Change: Audit Authentication Policy Change	**Success**
Policy Change: Audit Audit Policy Change	**Success and Failure**
Privilege Use: Audit Sensitive Privilege Use	**Success and Failure**
System: Audit System Integrity	**Success and Failure**
System: Audit Security System Extension	**Success and Failure**
System: Audit Security State Change	**Success and Failure**
System: Audit IPsec Driver	**Success and Failure**

Edit Audit Policies as follows:

Using Group Policy Management, access CorpNet.local's **Group Policy Objects** > **WorkgroupGPO.**

From Server Manager's menu bar, select **Tools** > **Group Policy Management** .

Maximize the windows for better viewing.

Expand **Forest: CorpNet.local** > **Domains** > **CorpNet.local** > **Group Policy Objects** .

Access the WorkstationGPO's Security Settings Local Policies.

Right-click **WorkstationGPO** and select **Edit** .

Maximize the windows for better viewing.

Under Computer Configuration, expand **Policies** > **Windows Settings** > **Security Settings > Local Policies** .

Modify Local Policies.

Select **Security Options** .

From the right pane, double-click the ***policy*** you want to edit.

Select **Define this policy setting** .

Select the ***policy settings*** as required.

Select **OK** .

Select **Yes** to confirm changes as necessary.

Repeat steps 3b - 3f for additional policy settings.

Modify the Event Log.

From the left pane, select **Event Log** .

From the right pane, double-click the ***policy*** you want to edit.

Select **Define this policy setting** .

Select the ***policy settings*** as required.

Select **OK** .

Modify Advanced Audit Policy Configuration.

From the left pane, expand **Advanced Audit Policy Configuration** > **Audit Policies** .

Select the audit policy category.

From the right pane, double-click the ***policy*** you want to edit.

Select **Configure the following audit events** .

Select the ***policy settings*** as required.

Select **OK** .

Repeat steps 5b–5f for additional policy settings.

12.3.5 Auditing Device Logs on a Switch (Demo Video)

Transcript:

In this demo, we'll discuss auditing device logs. and we'll use our Cisco managed switch as an example. Now, auditing device logs requires two things: First, we must enable logging so that we have log events that are being captured on the device. And second, we take the time to review or audit the logs to see what's happening. Let's see how this is done.

First, let's enable logging by clicking Administration, System Log, then Log Settings. You'll see that logging is currently enabled. This particular device has three types of logging. You can save log events in RAM, the random-access memory which will be lost after the system reboots. You can also save log events in flash memory which means that they'll be stored after the system reboots. And, finally, you can also aggregate, or compile, your logs and send them to a SYS log server.
To enable event aggregation, we enable the Syslog aggregator by checking the box and entering a time in seconds for Max Aggregation Time. Will enter 300 here, which will mean 5 minutes.
If you have a set of events that happen repeatedlyâ€”"for example, the same event keeps happening every five secondsâ€—"what will happen is the aggregator will say this event happened 15 times. Instead of seeing a separate entry for each similar event, you'll see the event happened and the number of times it happened within that aggregation period. We'll click apply to enable logging with aggregation.

You can add a remote log server, so you can view your logs from a central location. We'll click add and here you can specify the server by IP address or name, the UDP port, and have your logs sent to a certain facility. Let's click Close and go back to Log Settings.
As you'll notice there are several levels of logging on this device starting with Emergency. Let's talk about what each of these mean.
The Emergency level of severity means the system is not usable. An alert means that there's some sort of action that's needed. An event with critical severity means that the system is in a critical condition. An error means that there's an error in the system. Warning means that a system warning has occurred and a notice means that the system is functioning properly, but it just wants to let you know that something has happened. The informational level log is about device information. And then if you're having an issue with the device, you can log detailed debugging information about a particular event. As you'll notice we don't have debug information selected. But in our RAM logging , we're logging all of the other device events. Selecting a severity level causes all of those events higher to be selected as well. For example, we'll choose informational, and notice all the more critical events above it are selected, too.
We can also specify the Originator Identifier in our log messages by choosing None, or by hostname, IP address, or a description of your choice.
For flash memory logging which will be available after a system reboot, we're logging only a few top severity levels. These events are aggregating in our error log every 300 seconds, and then sent off to our defined SYS log server. Now let's take a look at some of the events in the event log.
With logging now setup and enabled, we're ready to view the logs. To see the logs, click on Status and Statistics, View Log, and then we can review both the RAM and Flash log entries. Notice there are various events by chronological order, newest on top. On this device, we see mostly notices and informational events, with an occasional warning. Keep an eye out for important messages like failed login attempts, logins from unauthorized workstations, device restarts, down links, and much more.
That's a quick summary of device log auditing. Remember to first enable logging so your device will generate events in the logs, then you need to take the time to periodically review the logs to see if there is important information there.

12.3.6 Enable Device Logs (Simulation)

Scenario

You are the IT security administrator for a small corporate network. You need to enable logging on the switch in the networking closet.

In this lab, your task is to:

Enable logging and the Syslog Aggregator.

Configure RAM Memory Logging as follows:

Emergency, Alert, and Critical: **Enable**

Error, Warning, Notice, Informational, and Debug: **Disable**

Configure Flash Memory Logging as follows:

Emergency and Alert: **Enable**

Critical, Error, Warning, Notice, Informational, and Debug: **Disable**

Copy the running configuration file to the startup configuration file using the following settings:

Source File Name: **Running configuration**

Destination File Name: **Startup configuration**

Explanation

Complete this lab as follows:

Access the Log Settings for the switch.

From the left menu, expand **Administration** > **System Log** .

Select **Log Settings** .

Enable Logging and Syslog Aggregator.

For Logging, select **Enable** .

For Syslog Aggregator, select **Enable** .

Configure RAM and Flash memory logging:

Under RAM Memory Logging *:*

Select **Emergency** , **Alert** , **and Critical** .

Clear **Error** , **Warning** , **Notice** , **Informational** , and **Debug** .

Under Flash Memory Logging:

Select **Emergency** and **Alert** .

Clear **Critical** , **Error** , **Warning** , **Notice** , **Informational** , and **Debug** .

Select **Apply** .

Save the changes.

From the top menu bar, select **Save** .

On the right, under *Source File Name* , make sure **Running configuration** is selected.

Under Copy/Save Configuration, select **Apply** .

Select **OK** .

Select **Done** .

12.3.7 Practice Questions (Section Quiz)

q_audits_attestation_secp8

You are a cybersecurity manager at a large multinational corporation. Your company has recently implemented a new security control system. You are tasked with ensuring the effectiveness and compliance of this new system.

Which of the following approaches would be the MOST effective for this task?

Answers:

Conduct an internal assessment only.

Conduct an external assessment only.

***Conduct both an internal and external assessment.**

Rely on self-assessment by the team that implemented the system.

Explanation:

Conducting both an internal and external assessment would be the most effective approach as it would provide a comprehensive evaluation of the new security control system. The internal assessment would align with the organization's needs and priorities, while the external assessment would provide an impartial and objective evaluation. This approach would also foster transparency and accountability, and facilitate knowledge sharing and professional development.

Conducting an internal assessment only would provide an in-depth evaluation of the organization's business processes. However, it may lack the objectivity and specialized expertise that an external assessment can provide. Therefore, this is not the most effective approach.

Conducting an external assessment only would provide an impartial and objective evaluation of the new security control system. However, it may not be as aligned with the organization's needs and priorities as an internal assessment would be. Therefore, this is not the most effective approach.

Relying on self-assessment by the team that implemented the system may not provide an objective evaluation of the new security control system. It may also lack the specialized expertise that an external assessment can provide. Therefore, this is not the most effective approach.

q_audits_auditing_secp8

Which of the following terms identifies the process of reviewing log files for suspicious activity and threshold compliance?

Answers:

Scanning

CompSec

***Auditing**

Phishing

Explanation:

Auditing is a complement to penetration testing and serves as documentation of attempted attacks that exceed preconfigured thresholds. Most operating systems, network devices, and security packages support the logging of usage data. Examples include

the success or failure of login attempts, file access, and administrative tasks. The detailed configuration of audit logs is necessary to ensure that all pertinent data is captured and available for review. Audit logs are sometimes used as evidence in court proceedings.

The following terms are not directly associated with the process of auditing:

Scanning is a term that describes digitizing an image, allowing it to be stored, modified, or understood by a computer.

CompSec is a company dedicated to solving complex IT problems.

Phishing is a type of social engineering attack often used to steal user data, including login credentials and credit card numbers.

q_audits_committee_secp8

A healthcare organization is developing its data privacy and security strategy. The leadership team is exploring different methods to monitor, evaluate, and improve security practices to ensure compliance with the Health Insurance Portability and Accountability Act (HIPAA).

What would be the MOST appropriate measure to maintain and oversee its privacy and security controls?

Answers:

***Establishing an audit committee**

Conducting a self-assessment

Implementing a security awareness program

Outsourcing security operations to a managed security service provider

Explanation:

Audit committees provide independent oversight and assurance regarding an organization's practices related to internal controls, risk management, and compliance.

Conducting a self-assessment can be a valuable tool for identifying strengths, weaknesses, and areas for improvement. Still, it provides a different level of independent oversight and continuity than an audit committee offers.

Implementing a security awareness program can help employees understand their role in maintaining data privacy and security. However, it does not provide the systematic oversight, review, and guidance on security practices that an audit committee offers.

Outsourcing security operations to a managed security service provider can help with the daily management and monitoring of security systems. However, the organization must still maintain its oversight and management capabilities.

q_audits_compliance_assessment_secp8

An organization must ensure its operating practices align with laws, regulations, standards, policies, and ethical requirements. The organization wants to evaluate the effectiveness of internal controls, identify any noncompliance or risk areas, and communicate findings to internal stakeholders such as risk managers.

Which internal assessment approach would be MOST appropriate for this purpose?

Answers:

Using an audit committee

***Compliance assessment**

Self-assessment

Regulatory assessment

Explanation:

A compliance assessment ensures that operating practices align with laws, regulations, standards, policies, and ethical requirements through an internal evaluation. This approach is the most appropriate for the scenario described.

While an audit committee may participate in the compliance process, it does not specifically evaluate compliance with laws and regulations.

A self-assessment, which can serve various purposes, including compliance, does not specifically evaluate compliance with laws and regulations.

External regulatory authorities or agencies typically perform a regulatory assessment to ensure compliance with specific laws, regulations, or industry standards, not as an internal assessment approach.

q_audits_external_assessment_secp8

An organization must hire nonaffiliated experts (consultants) to conduct an assessment. The organization expects the experts to provide a broad evaluation of their overall performance, practices, and capabilities, including specific focus areas, such as strategy, operational efficiency, risk management, cybersecurity, or compliance practices.

Which type of assessment meets the organization's needs?

Answers:

Self-assessment

Internal compliance assessment

***External assessment**

Internal audit committee

Explanation:

External experts or consultants conduct an external assessment to evaluate an organization's overall performance, practices, capabilities, or specific focus areas.

Individuals or organizations conduct self-assessments internally to evaluate their performance, practices, and adherence to established criteria against predetermined metrics and measures without the involvement of external experts.

An organization conducts an internal compliance assessment to ensure that its operating practices align with laws, regulations, standards, policies, and ethical requirements without involving external experts.

The internal audit committee provides independent oversight and assurance regarding an organization's financial reporting, internal controls, and risk management practices without involving external experts.

q_audits_external_examiniation_secp8

A multinational corporation wants to enhance its risk management procedures and validate that its systems, controls, and processes align with specific international standards, regulations, and best practices.

Which approach should the organization consider to ensure an unbiased and comprehensive analysis of its security posture?

Answers:

***External examination**

Self-assessment

Internal compliance assessment

Internal audit

Explanation:

External (unbiased) examination ensures a thorough assessment of the organization's compliance with international standards and best practices.

Self-assessment allows an organization to evaluate its performance, practices, and adherence to established criteria. While it can be helpful, self-assessment does not provide the independent and comprehensive review of an organization's security posture that an external examination would.

Internal compliance assessment, being an internal process, may not provide as comprehensive and unbiased an evaluation as an external examination.

An internal audit may offer a different level of objectivity and comprehensive assessment than an external examination can provide.

q_audits_external_secp8

Which type of audit is performed by either a consultant or an auditing firm employee?

Answers:

Internal audit

Usage audit

***External audit**

Financial audit

Explanation:

An external audit is performed by either a consultant or an auditing firm employee.

Internal audits are performed by an employee within an organization. They examine existing internal controls and maps the security structure for compliance with statutes and management goals.

Usage and financial audits can be performed by either internal or external auditors, depending on the reason for the audit.

q_audits_internal_01_secp8

Which of the following is true concerning internal audits?

Answers:

They are always highly rigorous.

The process is very formal.

***They are generally nonobjective.**

The auditor works independently.

Explanation:

Internal audits tend to be nonobjective and, consequently, may not be as rigorous.

None of the other answers best describe internal audits.

q_audits_internal_02_secp8

After several incidents, a financial organization is taking steps to improve its cybersecurity posture.

As part of these measures, it has implemented new security controls to ensure the confidentiality and integrity of customer data.

To provide the board of directors with reliable assurance about the effectiveness of these new controls before making a public statement, what approach should the organization employ?

Answers:

***Conduct an internal security audit.**

Initiate an attestation process.

Request an external security audit.

Commence third-party attestation.

Explanation:

Conducting an internal security audit can swiftly identify if the new measures are effective and directly communicate the findings to the board.

Initiating an attestation process offers a substantial degree of assurance, although it needs to address the board's need to understand the effectiveness of new measures from an internal perspective.

Requesting an external security audit may take more time and effort to evaluate the new measures than an internal audit. Additionally, the board requires assurance based on an internal evaluation.

Third-party attestation can offer high assurance, although it may offer a different, internally-focused assurance than the board seeks about the new security measures.

q_audits_pci_secp8

Which of the following standards relates to the use of credit cards?

Answers:

SOX

PoLP

***PCI DSS**

Financial audit

Explanation:

Personal Card Industry Data Security Standard (PCI DSS) compliance audits relate to the use of credit cards. These audits are regulated and enforced by major credit card companies.

A Sarbanes-Oxley (SOX) audit is a government audit by the SEC that relates to internal controls and focuses on IT security, access controls, data backup, change management, and physical security.

PoLP is the principle of least privilege. It does not apply in this scenario.

Financial audits are performed to ensure compliance with SOX or PCI DSS requirements.

q_audits_third_party_secp8

An organization validates its security controls, processes, and adherence to industry standards and wants an unbiased evaluation to instill confidence among stakeholders.

Which method should it employ for this purpose?

Answers:

Compliance assessment

Audit committee

Self-assessment

***Independent third-party audit**

Explanation:

An independent third-party audit offers an external, objective, and unbiased assessment of an organization's systems, controls, processes, and compliance. The goal is to instill confidence among stakeholders, including customers, business partners, regulatory bodies, and investors.

While compliance assessments help ensure that the organization's practices align with laws, regulations, standards, and policies, the company typically conducts them internally, making bias possible.

Although audit committees provide independent oversight and assurance regarding the organization's financial reporting, internal controls, and risk management practices, the organization's board members usually form the committees.

While self-assessments can be valuable for identifying improvement areas, they do not provide an external, unbiased perspective.

13.0 Data Protection and Compliance

13.1 Data Classification and Compliance

As you study this section, answer the following questions:

What is the purpose of classifying data?

What is the difference between PII and PHI?

Why is it important to know data destruction types?

How are privacy enhancing technologies used to protect PII?

Key terms for this section include the following:

Term	Definition
Data breach	When confidential or private data is read, copied, or changed without authorization. Data breach events may have notification and reporting requirements.
Escalated	In the context of support procedures, incident response, and breach-reporting, escalation is the process of involving expert and senior staff to assist in problem management.
Health Insurance Portability and Accountability Act (HIPAA)	US federal law that protects the storage, reading, modification, and transmission of personal healthcare data.
Regulated data	Information that has storage and handling compliance requirements defined by national and state legislation and/or industry regulations.
Trade secret	Intellectual property that gives a company a competitive advantage but hasn't been registered with a copyright, trademark, or patent.
Legal data	Documents and records that relate to matters of law, such as contracts, property, court cases, and regulatory filings.
Financial data	Data held about bank and investment accounts, plus information such as payroll and tax returns.
Human-readable data	Information stored in a file type that human beings can access and understand using basic viewer software, such as documents, images, video, and audio.

Non-human-readable data	Information stored in a file that human beings cannot read without a specialized processor to decode the binary or complex structure.
Data classification	The process of applying confidentiality and privacy labels to information.
Proprietary information or intellectual property (IP)	Information created by an organization, typically about the products or services that it makes or provides.
Data subjects	An individual that is identified by privacy data.
Data inventories	List of classified data or information stored or processed by a system.
Data retention	The process an organization uses to maintain the existence of and control over certain data in order to comply with business policies and/or applicable laws and regulations.
Disposal/decommissioning	In asset management, the policies and procedures that govern the removal of devices and software from production networks and their subsequent disposal through sale, donation, or as waste.
Sanitization	The process of thoroughly and completely removing data from a storage medium so that file remnants cannot be recovered.
Destruction	An asset disposal technique that ensures that data remnants are rendered physically inaccessible and irrevocable through degaussing, shredding, or incineration.
Certification	An asset disposal technique that relies on a third party to use sanitization or destruction methods for data remnant removal as well as providing documentary evidence that the process is complete and successful.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	3.3 Compare and contrast concepts and strategies to protect data. Data types Regulated

- Trade secret
- Intellectual property
- Legal information
- Financial information
- Human and non-human readable

Data classifications

- Sensitive
- Confidential
- Public
- Restricted
- Private
- Critical

General data considerations

- Data sovereignty

4.2 Explain the security implications of proper hardware, software, and data asset management.

Disposal/decommissioning

- Sanitization
- Destruction
- Certification
- Data retention

5.1 Summarize elements of effective security governance.

Roles and responsibilities for systems and data

- Controllers
- Processors

5.4 Summarize elements of effective security compliance.

	Consequences of non-compliance Fines Sanctions Reputational damage Loss of license Contractual impacts Privacy Legal implications Local/regional National Global Data subject Controller vs. processor Ownership Data inventory and retention Right to be forgotten

13.1.1 Consequences of Breaches (Lesson Video)

Transcript:

All too frequently, we read about a data breach. Large companies with security budgets in the millions of dollars are not exempt. Most of us have, at onetime or another, been notified that our information has been compromised. It could be bank accounts, credit cards, government records, or even health records. All these data breaches come with serious consequences. Let's talk about what they are.

A company's reputation determines if people will invest, if consumers will buy, and if foreign governments will allow a company to do business internationally. A company lives or dies by its revenues and investments.
A breach that exposes private client data directly affects the way consumers and investors spend money. A data breach can cause stock prices to fall. Falling stock prices lead to sell offs and can permanently damage a company.
When trust is broken and consumer confidence is damaged, revenues drop because consumers stop using the company. Also, the loss of government contracts because of breach can prove fatal for a company.
Twenty years ago, the term identity theft didn't have any real meaning to the everyday person. Fast forward to today, and it is difficult to find someone who hasn't experienced identity theft on some level. When a breach occurs and personal

information is stolen, the affected person is forced to do hours, days, and even weeks of work to correct someone else's mistake. The effects of identity theft can last years and cause fear and paranoia among those affected. Lives have been turned upside down leaving victims frustrated because it is almost impossible to catch the perpetrator.

A data breach can also be financially costly. With the introduction of GDPR and the California Consumer Privacy Act the number of companies that can be fined has grown significantly. In the United States, HIPAA has always carried fines for companies found in non-compliance. Companies that profit from the use of consumer personal information can also have heavy fines levied by the Federal Trade Commission and possibly the Security and Exchange Commission.

To give you an idea of the magnitude of these fines, in 2019 three data breaches at three separate companies cost the companies over 1 billion dollars. It is estimated that fines for a single data breach now averages $6.45 million dollars.

Intellectual Property is the lifeblood of companies. When IP is stolen ,companies lose competitive advantage.

The internet has made the world smaller. Companies are now competing with others from around the globe. IP stolen through a data breach is often sold to competing companies.

Many of the overseas companies are not subject to laws in the United States. This makes it difficult to enforce IP laws. The reverse engineering or the direct copy of IP gives thieves an undeserved revenue source. IP theft also floods the market with counterfeit goods.

Let's look at an IP theft scenario. The XYZ company has worked tirelessly on a new product that's cutting-edge. As the company is going into production, they experience a data breach. They are not sure what data may have been lost. A few months later, counterfeit versions of their cutting-edge product start flooding the market. They now discover that their intellectual property was stolen during the breach. The fakes are selling for less money and destroying XYZ's profits.

Escalation can be separated into two categories, internal escalation and external escalation. Internal escalation is part of a company's incident response plan. This usually means a tiered approach to informing management of a breach. As a breach is investigated, the findings and severity will determine how far the escalation will go. The external escalation involves experts brought in from the outside to investigate, provide legal counsel, or even prosecute, which involves law enforcement. The key is to know when a situation has surpassed your abilities to resolve it.

That's it for this lesson. In this lesson, we showed you how data breaches can have serious consequences including loss of trust, fines, and theft of intellectual property. Fines are averaging $6.45 million per event. All 50 states have laws outlining the data breach notification requirements. We also looked at a scenario of IP theft and its consequences.

13.1.2 Consequences of Breaches Facts

Breaches affect all companies. This includes even large companies with security budgets in the millions of dollars. Most of us have, at one time or another, been notified that our information has been compromised. It could be bank accounts, credit cards, government records, or even health records. All these data breaches come with serious consequences.

This lesson covers the following topics:

- Impacts of non-compliance
- Privacy breaches and data breaches
- Software licensing non-compliance
- Impacts of contractual non-compliance

Impacts of Non-compliance

Security compliance refers to organizations' adherence to applicable security standards, regulations, and best practices to protect sensitive information, mitigate risks, and ensure data confidentiality, integrity, and availability. Effective compliance necessitates establishing and implementing policies, procedures, controls, and technical measures to meet the requirements set forth by regulatory bodies, industry standards, and legal obligations.

Non-compliance with data protection laws and regulations can have severe consequences for organizations. The consequences vary depending on jurisdiction and the specific regulations violated. Common ramifications for non-compliance include legal sanctions such as financial penalties, legal liabilities, reputational damage, and loss of customer trust. Sanctions refer to penalties,

disciplinary actions, or measures imposed due to non-compliance with laws, regulations, or rules. Sanctions are enforced by governing bodies, regulatory authorities, or organizations overseeing the specific domain in which the non-compliance occurred. Regulatory agencies may impose substantial fines, which can amount to millions or even billions of dollars, depending on the severity of the violation. Legal action from affected individuals or data subjects may lead to costly lawsuits and settlements. Non-compliance can harm an organization's reputation, eroding customer trust, decreasing business opportunities, and potentially losing contracts or partnerships. Organizations may also face additional regulatory scrutiny, including increased audits, investigations, or mandated remediation measures. Organizations must prioritize data protection compliance, implement appropriate security measures, conduct regular risk assessments, and stay informed about evolving data protection laws and regulations to avoid these consequences.

Privacy Breaches and Data Breaches

A data breach occurs when information is read, modified, or deleted without authorization. "Read" in this sense can mean either seen by a person or transferred to a network or storage media. A data breach is the loss of any type of data (but notably corporate information and intellectual property). In contrast, a privacy breach refers specifically to the loss or disclosure of personal and sensitive data.

Organizational Consequences

A data or privacy breach can have severe organizational consequences:

Consequences	Description
Reputation damage	Data breaches cause widespread negative publicity, and customers are less likely to trust a company that cannot secure its information assets. Consumers stop using the company, revenues drop, and possibly lost investments.
Identity theft	If the breached data is exploited to perform identity theft, the data subject may be able to sue for damages. The effects of identity theft can last for years and cause fear and paranoia among the victims. Lives have been turned upside down, leaving victims frustrated to deal with the ramifications.
Fines	Legislation might empower a regulator to levy fines. These can be a fixed sum or, in the most serious cases, a percentage of turnover. Regulations like the General Data Protection Regulation (GDPR), California Consumer Privacy Act (CCPA), Health Insurance Portability and Accountability Act (HIPAA), and many more do impose fines for non-compliance.
IP theft	Loss of company data can lead to loss of revenue. This typically occurs when copyright material—unreleased movies and music tracks—is breached. The loss of patents, designs, trade secrets, and so on to competitors or state actors can also cause commercial losses, especially in overseas markets where IP theft may be difficult to remedy through legal action.

Escalation

A breach may be detected by technical staff, and if the event is considered minor, there may be a temptation to remediate the system and take no further notification action. This could place the company in legal jeopardy. Any breach of personal data and most breaches of IP should be escalated to senior decision-makers, and any impacts from legislation and regulation properly considered.

Notifications of Breaches

The requirements for different types of breaches are set out in law and in regulations. The requirements indicate who must be notified. A data breach can mean the loss or theft of information, the accidental disclosure of information, or the loss or damage of information. Note that there are substantial risks from accidental breaches if effective procedures are not in place. If a database administrator can run a query that shows unredacted credit card numbers, that is a data breach, regardless of whether the query ever leaves the database server.

Depending on the regulations, a breach may be considered to have occurred if there is just the potential for unauthorized access. For example, if a personal data file is configured with permissions that mistakenly allow any authenticated user to read it, this could be classified as a notifiable data breach, even if audit logs show that no improper access attempts were made.

Public Notification and Disclosure

Other than the regulator, notification might need to be made to law enforcement, individuals, and third-party companies affected by the breach and the public through press or social media channels. For example, the Health Insurance Portability and Accountability Act (HIPAA) sets out reporting requirements in legislation, requiring breach notification to the affected individuals, the Secretary of the US Department of Health and Human Services, and, if more than 500 individuals are affected, to the media (hhs.gov/hipaa/for-professionals/breach-notification/index.html). The requirements also set out timescales for when these parties should be notified. For example, under GDPR, the notification must be made within 72 hours of becoming aware of a breach of personal data (csoonline.com/article/3383244/how-to-report-a-data-breach-under-gdpr.html). Regulations will also set out disclosing requirements or the information that must be provided to each of the affected parties. Disclosure is likely to include a description of what information was breached, details for the main point of contact, likely consequences arising from the breach, and measures taken to mitigate the breach.

GDPR offers stronger protections than most federal and state laws in the United States, which tend to focus on industry-specific regulations, narrower definitions of personal data, and fewer rights and protections for data subjects. The passage of the California Consumer Privacy Act (CCPA) has changed the picture for domestic US legislation, however (csoonline.com/article/3292578/california-consumer-privacy-act-what-you-need-to-know-to-be-compliant.html).

Software Licensing

Non-compliance with software licensing requirements can result in the revocation of usage rights and other consequences, such as fines. Violations of license agreements, such as exceeding permitted installations, unauthorized sharing, or other unauthorized usage, constitute contractual non-compliance. Other forms of non-compliance include breaching license terms, such as modifying code or distributing software without authorization. In response, software vendors or licensing authorities may revoke or suspend licenses and take other legal actions. The loss of software licenses can disrupt business operations, causing inefficiencies and workflow interruptions as well as cause significant reputational damage. To ensure compliance, organizations can rectify non-compliance through license remediation, proper license management, and audits.

Impacts of Contractual Non-compliance

Category	Description
Breach of contract	Non-compliance can result in a breach of contract. Contracts between parties often include provisions related to data protection, cybersecurity measures, and the safeguarding of sensitive information. Failure to meet these contractual obligations can lead to legal consequences, including potential liability for damages or loss the non-compliant party suffers.
Termination of contracts	Non-compliance may give the non-compliant party grounds for contract termination. Contractual agreements may contain clauses allowing termination if the other party fails to adequately protect data or implement sufficient cybersecurity measures . The non-compliant party may face termination penalties, loss of business relationships, and the need to seek new contractual arrangements, which will be complicated by poor past performance.

Category	Description
Indemnification and liability	Non-compliance may result in the non-compliant party assuming liability for damages caused by a security breach or data loss. Contractual agreements may include indemnification clauses, shifting responsibility for losses or legal expenses resulting from cybersecurity incidents onto the non-compliant party, leading to financial burdens and reputational damage.
Non-compliance penalties	Contracts may stipulate penalties or financial consequences for non-compliance with cybersecurity requirements, such as monetary fines or contractual damages that the n on-compliant party must pay to the aggrieved party. Non-compliance penalties aim to incentivize adherence to cybersecurity measures outlined in the contractual agreement.

13.1.3 Information Classification (Lesson Video)

Transcript:

In this video, we're going to talk about information classification. There are several types of data, and data needs to be classified and secured. First, let's go over general classifications of data.

All data falls into one of five categories. Public data is available to anyone. Private data has limited distribution. Sensitive data is for internal use. Critical data is covered by laws and compliance regulations. Proprietary data is exclusive to a particular entity. Let's look at each of these in detail.
The first information classification is Public. This classification allows anyone to have a copy of the information. Examples include a public website that everyone can access, company brochures, and marketing material. This data exists to be released.
Private data is not for general release to the public. It includes internal email addresses, company meeting notes, company presentations, and customer lists. This is information you don't want released to the public, but if it were accidentally released, damage to your organization would be light to moderate.
Sensitive data is information that, if exposed, could cause significant damage to an organization. Sensitive data is restricted even within the company that's responsible for it. It may include strategic information, customer lists, personally identifiable information (PII), and company strategies.
Critical data includes personally identifiable information, credit card data, health records, and other information covered under compliance regulations. These include The General Data Protection Regulation, or GDPR, and The California Consumer Protection Act. The disclosure of critical data will cause severe harm and financial penalties.
Proprietary information includes information that's unique and differentiates a company from its competitors. Examples are Google's search algorithm, the recipe for Coca-Cola, or patents granted for inventions and processes. The release of this information would cause serious damage, maybe even irreparable damage.
Now let's talk about very specific kinds of data that need to be classified by a company. This data is covered by laws, government regulations, and compliance acts.
The first is classified data.
Personal Identifiable Information includes almost any data that can be used to identify a person: their name, addresses, email, phone number, or even a photograph.
The Health Insurance Portability and Accountability Act (HIPAA) has many facets, but the one that concerns you most as a security administrator is their policy on Electronic Health Information, or eHI. This classification includes data related to an individual's health and information collected by doctors, health clinics, health insurance companies, and employers. It must be secured in all usage statesâ€“ at rest, in use, or in motion. The exposure, loss, or breach of this data incures significant fines from the U.S. government.
Financial data is covered by a few laws and compliance bodies. It includes consumer credit card information, bank account numbers, and corporate financial requirements.
There are also government/military classification levels. The public can view Unclassified information. The next classification is Sensitive, but Unclassified. For this classification, the government or military would prefer the information not be available to the public, but it is unclassified. Disclosure could cause harm, but the harm would be minimal.

The next level is Confidential. It allows restricted information release under the Freedom of Information Act. The level above that is Secret, and it includes troop movements, deployments, and overall capabilities. The highest classification is Top Secret. Release of Top-Secret information can pose a grave threat to the country and national security.
That's it for this lesson. In this lesson, we covered the different data classifications and examples of data in each classification: public, private, sensitive, and critical.
Then we looked at specific types of data, including PII, ePHI, credit cards, and unclassified, confidential, secret, and top-secret.

13.1.4 Information Classification Facts

The concept of data types refers to categorizing or classifying data based on its inherent characteristics, structure, and intended use. Data types provide a way to organize and understand the different data forms within a system or dataset. Classifying data into specific types makes analyzing, processing, interpreting, and securing information easier.

This lesson covers the following topics:

- Regulated data
- Trade secrets
- Legal and financial data
- Human-readable and non-human-readable data
- Data classifications

Regulated Data

Regulated data refers to specific categories of information subject to legal or regulatory requirements regarding their handling, storage, and protection. Regulated data typically includes sensitive or personally identifiable information (PII) protected by laws and regulations to ensure privacy, security, and appropriate use. The types of regulated data vary depending on jurisdiction and the specific regulations applicable to the organization or data. Common examples of regulated data include financial information, healthcare records, social security numbers, credit card details, and other personally identifiable information. Privacy laws and industry-specific regulations often protect these data types, such as the Health Insurance Portability and Accountability Act (HIPAA) for healthcare data or the Payment Card Industry Data Security Standard (PCI DSS) for credit card information. Organizations that handle regulated data must comply with relevant laws and regulations governing its protection. Compliance typically involves implementing appropriate security measures, data encryption, access controls, data breach notification procedures, and data handling protocols. Organizations may also need to establish data storage, retention, and destruction safeguards to meet regulatory requirements.

Trade Secrets

Trade secret data refers to valuable, confidential information that gives a business a competitive advantage. Trade secrets encompass much nonpublic, proprietary information, including formulas, processes, methods, techniques, customer lists, pricing information, marketing strategies, and other business-critical data. Trade secrets have commercial value derived from their secrecy. Businesses often require employees and contractors to sign non-disclosure agreements (NDAs) to safeguard the confidentiality of trade secrets. Disclosure or unauthorized use of trade secret data is a serious legal matter. Companies can take legal action against individuals or organizations unlawfully acquiring, using, or disclosing trade secrets. Laws related to trade secrets vary across jurisdictions, but they generally aim to prevent unfair competition and provide remedies for misappropriation.

Legal and Financial Data

Legal and financial data encompass critical data for legal compliance, financial reporting, decision-making, and risk management. Legal data includes documents, contracts, legal agreements, court records, litigation information, intellectual property filings, regulatory filings, and other legal documents. It may also encompass information related to corporate governance, compliance with laws and regulations, and legal obligations specific to an industry or jurisdiction. On the other hand, financial data pertains to information concerning an organization's financial activities, performance, and transactions, including financial statements, balance sheets, income statements, cash flow statements, audit reports, tax records, financial projections, budgets, and other financial reports. Financial data also encompasses details of financial transactions, such as accounts payable, accounts receivable, general ledger entries, and transactional records. Legal and financial data are highly sensitive and confidential due to their nature and the potential impact they can have on an organization's reputation, legal standing, and financial stability.

Human-Readable and Non-Human-Readable Data

Human-readable data refers to information humans can easily understand and interpret without additional processing or translation. Human-readable data describes a format that is accessible and readable, such as text, images, or multimedia content. Examples of human-readable data include documents, reports, emails, web pages, and presentations. On the other hand, non-human-readable data refers to data that is not easily understood or interpreted by humans in its raw form. It may be in a machine-readable format, such as binary code, encrypted data, or data represented in a complex structure or encoding that requires specialized software or algorithms to decipher and interpret. Non-human-readable data often requires additional processing or transformation to make it understandable to humans.

Human-readable and non-human-readable data formats have distinct implications for security operations and controls. Human-readable and non-human-readable data formats impact security operations and controls in different ways. Security monitoring, user awareness, DLP, content filtering, and web security are more directly applicable to human-readable data formats.

On the other hand, encryption, access controls, intrusion detection and prevention, secure data exchange, and code/application security are more relevant to non-human-readable data formats. It is important to note that non-human-readable data formats can impede the capabilities of security controls because non-human-readable data formats cannot be easily interpreted using traditional methods and require specialized approaches to inspect and protect them. A comprehensive security approach considers both types of data formats and implements appropriate measures to protect them based on their characteristics and associated risks.

Data Classifications

Data classification and typing schemas tag data assets so that they can be managed through the information lifecycle. A data classification schema is a decision tree for applying one or more tags or labels to each data asset. Many data classification schemas are based on the degree of confidentiality required:

Public (unclassified) — there are no restrictions on viewing the data. Public information presents no risk to an organization if it is disclosed but does present a risk if it is modified or unavailable.

Confidential (secret) — the information is highly sensitive, for viewing only by approved persons within the owner organization and possibly by trusted third parties under NDA.

Critical (top secret) — the information is too valuable to allow any risk of its capture. Viewing is severely restricted.

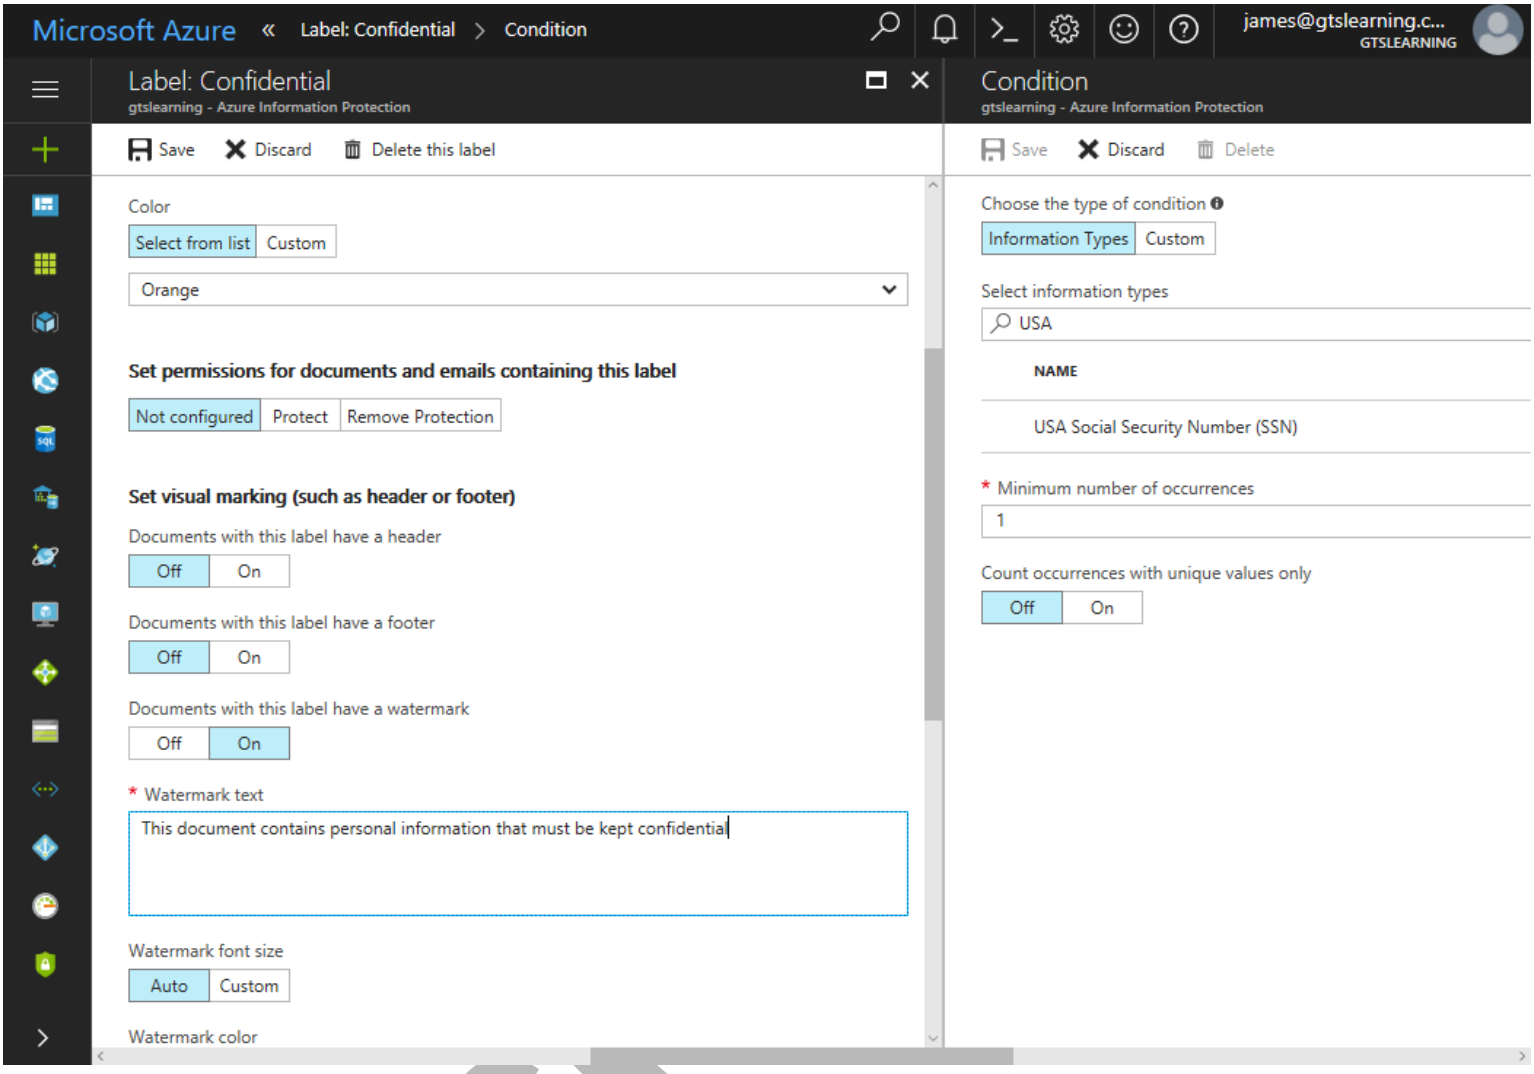

Using Microsoft Azure Information Protection to define an automatic document labeling and watermarking policy. (Screenshot used with permission from Microsoft.)

Another type of classification schema identifies the kind of information asset:

Classification	Description
Proprietary	Proprietary information or intellectual property (IP) is information created and owned by the company, typically about the products or services they make or perform. IP is an obvious target for a company's competitors, and IP in some industries (such as defense or energy) is of interest to foreign governments. IP may also represent a counterfeiting opportunity (movies, music, and books, for instance).
Private/personal data	This information relates to an individual identity. Private data examples include personally identifiable information (PII) such as names, addresses, social security numbers, financial information, and sensitive data like health records, login credentials, biometric data, and confidential business information.

Classification	Description
Sensitive	This label is usually used in the context of personal data privacy-sensitive information about a subject that could harm them if made public and could prejudice decisions made about them if referred to by internal procedures. As defined by the EU's General Data Protection Regulation (GDPR), sensitive personal data includes religious beliefs, political opinions, trade union membership, gender, sexual orientation, racial or ethnic origin, genetic data, and health information (ec.europa.eu/info/law/law-topic/data-protection/reform/rules-business-and-organisations/legal-grounds-processing-data/sensitive-data/what-personal-data-considered-sensitive_en).
Restricted	This classification refers to sensitive information that requires stringent controls and limited access due to its highly confidential nature. Restricted data typically includes data that, if disclosed or accessed by unauthorized individuals, could cause significant harm to individuals, organizations, or national security.

Government and Military classifications can correspond to security-clearance levels used for military members, government officials, and civilian contractors. The structure will be similar to the following:

Government Classification	Explanation
Unclassified	Data that is available for public use. No clearance is needed.
Sensitive	Sensitive data is unclassified, but the government would prefer that the data not be public.
Confidential	Confidential data is information with limited distribution. No clearance is required.
Secret	The first data classification that requires a security clearance.
Top Secret	The highest level of data classification. Only those with Top Secret clearance can view the documents on a need-to-know basis.
Top Secret Compartmentalized	This data is not any more secret than Top Secret but has very specific and limited distribution.

13.1.5 Privacy and Responsibility of Data (Lesson Video)

Transcript:

We're going to spend a few minutes reviewing some of the laws that have been enacted to protect the privacy of electronic data.

The United States has passed a significant number of laws regarding information security, especially the privacy of electronic data. Prior to the year 2000, there weren't many laws that addressed this particular area of information security. However, due to a rash of break-ins, exploits, and another abuses, many states and the federal government enacted laws to protect the privacy of electronic data.
A key thing to remember is that you need to be very familiar with the requirements of the laws that apply to your organization. To protect your organization and to protect yourself, you need to be aware of how these laws apply to your organization.

This lesson addresses some of the higher profile data privacy laws, but it doesn't identify all the laws that apply to your organization, nor does it provide minute details of each law that it does cover. This lesson provides only a simple, high-level overview of these laws. You are responsible to research whether these laws apply to your organization, and if so, how. You are also responsible to be aware of other federal, state, and even local laws that effect your organization.

Let's begin with the Health Insurance Portability and Accountability Act of 1996 called HIPAA. HIPAA specifies that all organizations having anything to do with healthcare must protect the health information that they maintain.

These organizations must implement policies and procedures to protect this information, regardless of the storage medium, such as paper in a filing cabinet or electronic digital format. HIPAA also establishes national standards for transferring electronic healthcare information.

The next act you need to be familiar with is the Sarbanes-Oxley Act of 2002. The Sarbanes-Oxley Act of 2002, sometimes referred to as Sarbox, is the result of a flood of corporate fraud in the late 1990s and the early 2000s. Sarbox requires publically traded companies to adhere to very stringent reporting requirements and implement strong controls on electronic financial reporting systems.

A key point in Sarbox is that organizations have to keep information for a certain time. This especially relates to email. In fact, Sarbox created an entire industry of email archiving companies so that organizations can comply with its regulations.

Next, we need to look at the Gramm-Leach-Bliley Act, which we refer to as GLBA. GLBA is designed to protect private data much like HIPAA does; however, GLBA applies to private information held at financial institutions.

There are two main functions of GLBA. First, it requires all banks and any other financial institutions to alert their customers as to that organization's privacy policies. If you've received a little pamphlet from your financial institution that contained that institution's privacy statement, that pamphlet was in response to the Gramm-Leach-Bliley Act.

In addition, all personally identifiable financial information (PII) within a financial institutionâ€”either electronic or paper formatsâ€”has to be protected. Essentially this act specifies that a policy must be in place to protect private information from foreseeable threats and to maintain data integrity. In order to implement this, GLBA requires financial institutions to put in place three main components.

The first one is the Financial Privacy Rule. The second one is the Safeguards Rule. And the third one is called Pretexting Protection. Essentially, the Financial Privacy Rule requires all financial institutions to provide each customer with the privacy notice that we talked about earlier. It must be provided at the time the relationship is established. That is, when you go in to open a new banking account. It also has to be provided every year thereafter.

The Safeguards Rule requires financial institutions to develop a written information security plan, which describes in detail how the company plans to protect clients' personal information. Finally, Pretexting Protection encourages financial institutions to train their staff how to recognize social engineering exploits, which they call pretexting. The reason is that social engineering exploits use some type of pretext.

You should also be familiar with the Patriot Act of 2001. This act enables law enforcement agencies to detect and suppress terrorism by giving law enforcement the authority to request information from organizations. All organizations, public or private, must provide the requested information to the appropriate law enforcement agencies under the authority of a valid court order or a subpoena.

The next law you need to be familiar with is the California Database Security Breach Act of 2003. This law specifies that any agency, person, government entity, or company that does business in the state of California must inform California residents within 48 hours if a database breach or other security breach occurs in which personal information has been stolen or is believed to have been stolen.

Most other states have similar state laws modeled on the California Database Security Breach Act of 2003, making it a significant act. Therefore, use your favorite search engine to search for similar acts in the states in which you do business. If you do business in those states and your organization has a breach of personal information, then you are under a regulatory compliance requirement to inform anyone who may be affected.

Next, let's look at the Children's Online Privacy Protection Act of 1998, which is referred to as COPPA.

COPPA requires organizations that provide online services designed for kids below the age of 13â€”such as websites and gaming sitesâ€”to obtain parental consent prior to collecting a child's personal information and using it, such as displaying it on the website, selling it to a marketing company, and so on.

In addition, COPPA also specifies that such online services cannot require kids to provide more information than what they determine as reasonable in order to participate.

The General Data Protection Regulation, or GDPR, is a data compliance regulation that started in 2018. GDPR provides sweeping changes to the way customer data is treated in the European Union. The regulation not only applies to countries that are part of the EU, but it also applies to any company that has customers that reside in the EU itself. The regulation requires companies to be prepared to tell costumers what data they have, where they got it, who they've given it to, and how it's being used. It also gives customers the right to be forgotten in some cases. Failure to adhere to the regulations can be costly.

The last one we'll look at is the California Consumer Privacy Act, or CCPA, was passed in 2020 and was one of the first data privacy acts in the United States. The act applies to California state residents, but companies that do business in the state and

meet other benchmarks must also comply. The consumers' rights include the right to know what data is collected and whether their data was sold and to whom. The consumer is also allowed to access their own data and able to deny the sale of it. Under certain circumstances, they're allowed to have data deleted and cannot be discriminated against for doing so.
That's it for this video. This is not an all-inclusiveÂ list of laws and acts regarding electronic data privacy with which you need to be familiar. GDPR and CCPA are both relatively new, so make sure you research laws that apply to your organization beyond just these two. It's also probably a good idea to use the servicesÂ of a lawyer who's familiar with this field just toÂ make sure that your organization stays in compliance.Â Please note again that non-compliance to these laws and regulations can be very costly.

13.1.6 Privacy and Responsibility of Data

Privacy data refers to personally identifiable or sensitive information associated with an individual's personal, financial, or social identity, including data that, if exposed or mishandled, could infringe upon an individual's privacy rights.

This lesson covers the following topics:

- Privacy data
- Legal implications
- Roles and responsibilities
- Right to be forgotten
- Ownership of privacy data
- Data inventories and retention

Privacy Data

Examples of privacy data include names, addresses, contact information, social security numbers, medical records, financial transactions, and, generally, any other data that can be used to identify a specific person. Privacy data and confidential data have certain similarities. Both types of data require protection due to their sensitive nature. Unauthorized access, disclosure, or misuse of privacy or confidential data can negatively affect individuals or organizations.

Additionally, privacy and confidential data are subject to legal and ethical considerations. Organizations must comply with relevant laws and regulations, such as data protection and privacy laws, to safeguard both data types. However, there are also notable differences between privacy data and confidential data.

Confidential data encompasses any information that requires protection due to its confidential nature, regardless of whether it pertains to an individual. Examples include trade secrets, intellectual property, financial statements, proprietary algorithms, source code, and other nonpublic information. Privacy data, on the other hand, specifically refers to information that can identify or impact an individual's privacy. Confidential data is primarily concerned with safeguarding information from unauthorized access, use, or disclosure to maintain business competitiveness, protect intellectual property, or preserve the integrity of sensitive company data.

Privacy data focuses on protecting personal information to preserve an individual's privacy rights, prevent identity theft, and maintain the confidentiality of personal details. Privacy data is closely associated with the rights of individuals to control the use and disclosure of their personal information. Individuals have the right to access, correct, and request the deletion of their privacy data. In contrast, confidential data typically does not grant specific rights to the data subjects , as it relates more to organizations' proprietary information. The handling of privacy data often requires explicit consent from the data subject for its collection, use, and disclosure, particularly in compliance with privacy laws and regulations. On the other hand, confidential data, while protected, may not necessarily require individual consent for its handling, as it is associated with internal or business-related information.

Privacy and confidential data share similarities in sensitivity and legal considerations. However, scope, focus, data subject rights, and consent requirements differ. While both types of data require careful handling and protection, privacy data pertains explicitly to personal information and individual privacy rights.

Legal Implications

Protecting privacy data carries significant local, national, and global legal implications. Many countries have specific privacy laws and regulations that dictate how personal data should be handled within their jurisdiction. These laws define the rights of individuals, the responsibilities of organizations, and the procedures for data protection and privacy enforcement. At the national level, data protection authorities or supervisory bodies enforce privacy laws and oversee compliance. They have the authority to investigate data breaches, issue fines, and take legal action against organizations that fail to protect privacy data or violate individuals' privacy rights. The General Data Protection Regulation (GDPR) in the European Union has had a substantial impact globally by setting high privacy and data protection standards. GDPR applies to organizations that process the personal data of EU residents, regardless of their physical location. This extraterritorial effect ensures that organizations worldwide adhere to GDPR principles when handling EU citizens' personal data. Cross-border data transfers are also subject to specific requirements and restrictions. For example, the GDPR restricts transferring personal data outside the European Economic Area unless adequate safeguards exist to protect privacy data. Understanding and adhering to these legal requirements are essential to avoid legal consequences, maintain trust with individuals, and foster a global culture of privacy and data protection.

Roles and Responsibilities

Data controller and data processor are two distinct roles defined under data protection regulations, such as the General Data Protection Regulation (GDPR). Although they deal with personal data, these roles have important similarities and differences. The data controller and data processor are involved in handling personal data. Both roles are responsible for ensuring personal data protection in compliance with data protection laws and regulations. The data controller and data processor must also adhere to data protection laws. They are required to process personal data lawfully, securely, and transparently.

The primary distinction lies in their roles and responsibilities. The **data controller** is the entity or organization that determines the purposes and means of processing personal data. They have overall control and responsibility for the processing of personal data. The data controller decides why and how personal data is processed. They exercise decision-making authority, define the purposes of data processing, and determine the categories of data to be processed. Data controllers have direct legal obligations and responsibilities under data protection laws. They are accountable for handling compliance, obtaining appropriate consent from data subjects, providing privacy notices, implementing data protection policies and procedures, and handling data subject requests. The **data processor** processes personal data on behalf of the data controller. They act under the authority and instructions of the data controller. Data processors do not have independent decision-making power over personal data. They process data solely as instructed by the data controller. Data processors have legal obligations to process personal data only for the purposes defined by the data controller. They must implement appropriate security measures, maintain the confidentiality and integrity of the data, and cooperate with the data controller to meet their legal obligations. Data processors are also required to keep records of their processing activities. Data processors include cloud service providers or payroll processing companies, for example.

A data subject refers to an individual whose personal data is processed by an organization or other entity. They are the individuals to whom the personal data refers. Data subjects hold certain rights and protections under data protection laws, such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). One of the rights afforded to data subjects is the right of access, meaning that data subjects have the right to request access to their personal data and obtain information about how it is being processed. Subjects can inquire about the purposes of processing, the categories of data being processed, recipients of the data, and the duration of data retention. Data subjects have the right to rectification, which means that if data subjects discover that the personal data held by an organization is inaccurate or incomplete, they have the right to request its correction to ensure that their personal data is up to date and accurate. Data subjects also have the right to request the erasure or removal of their personal data under certain circumstances.

For example, if the data is no longer necessary for the purposes it was collected, or if the data subject withdraws their consent for its processing, they can request its deletion. Data subjects can request the restriction of processing their personal data. The implications are that while privacy data can still be stored, it cannot be processed further except under specific conditions. This

right gives data subjects control over their personal data's ongoing use. Data portability is another right granted to data subjects. Subjects have the right to receive their personal data in a commonly used and machine-readable format, ensuring their ability to move and transfer their personal information as desired. Data subjects have the right to object to processing their personal data based on specific grounds. Examples include if a subject believes their data is being processed for purposes that are not legitimate or if they wish to object to direct marketing activities. Lastly, data subjects have the right to withdraw their consent for the processing of their personal data. If the processing is based on their consent, they can revoke it at any time, and the organization must cease processing the data accordingly.

Data subjects exercise these rights by contacting the data controller, who ensures that data subject rights are respected, facilitating the exercise of these rights, and addressing any concerns or requests from data subjects.

Right to Be Forgotten

The "right to be forgotten" is a fundamental principle outlined in the General Data Protection Regulation (GDPR) that grants data subjects the right to request the erasure or deletion of their personal data under certain circumstances. It empowers individuals to have their personal information removed from online platforms, databases, or any other sources where their data is being processed and made publicly available.

The right to be forgotten recognizes the importance of individual privacy and control over personal data. Upon receiving a valid erasure request, the data controller must erase the personal data promptly unless there are legitimate grounds for refusing the request. This right extends to the removal of data from the organization's systems and to any third parties with whom the data has been shared or made publicly available. This right may be limited if the processing of personal data is necessary for exercising the right of freedom of expression and information, compliance with a legal obligation, or the establishment, exercise, or defense of legal claims. The right to be forgotten serves as a mechanism for individuals to regain control over their personal information. It promotes privacy and data protection by enabling subjects to remove personal data when it is no longer necessary or lawful to retain it.

Ownership of Privacy Data

The question of ownership regarding privacy data is a complex topic. In general, it is not easy to attribute traditional notions of ownership to privacy data. The ownership of privacy data is often not considered in terms of traditional property rights. Under many data protection laws, such as the GDPR, the emphasis is placed on the rights and protections of the data subject rather than determining ownership. The data subject has control over their personal data and can exercise certain rights, such as the right to access, rectify, and delete their data.

However, organizations that collect and process personal data are considered custodians or stewards of the data rather than owners. They have legal and ethical responsibilities to handle personal data securely and lawfully and to respect the rights of the data subjects. It is important to note that privacy data often consists of information about individuals, and those individuals have a strong interest in protecting their personal information. Data protection laws aim to provide individuals with control and protection over their personal data, ensuring transparency, consent, and fair processing practices. While the concept of ownership might not directly apply to privacy data, individuals have rights and control over their personal information, and organizations are legally accountable for handling the data responsibly. The focus is on safeguarding privacy rights and ensuring data protection rather than assigning ownership in the traditional sense.

Data Inventories and Retention

Privacy laws profoundly impact data inventories and data retention practices within organizations. These laws, such as the GDPR and CCPA, require organizations to maintain a detailed record of the personal data they collect, process, and store. Data inventories provide a comprehensive overview of the types of data being handled, the purposes for processing, the legal basis, and recipients of the data to ensure transparency and accountability, as organizations can clearly understand and document their data processing activities in compliance with privacy laws. Privacy laws stipulate that organizations must have a lawful basis for processing personal data. Data inventories are crucial in identifying the legal grounds for data processing. By documenting the legal basis for each category of personal data, organizations can ensure that their processing activities align with the specified lawful purposes outlined in privacy laws. Organizations must collect and process only the necessary elements of personal data for

specific and legitimate purposes. Data inventories assist organizations in evaluating the personal data they collect, ensuring they only gather necessary information. By keeping the inventory up to date, organizations can align their practices with the principles of data minimization and purpose limitation.

Data retention is another area impacted by privacy laws. Organizations must retain personal data only for as long as necessary to fulfill the intended purpose or as required by law. Data inventories help organizations determine appropriate retention periods for different categories of personal data, ensuring compliance with data storage limitation requirements. Keeping accurate records in the data inventory enables organizations to securely delete or anonymize data when it is no longer needed. Privacy laws grant individuals various rights, such as the right to access their personal data. Data inventories are instrumental in facilitating the exercise of these rights. By maintaining comprehensive inventories, organizations can promptly respond to data subject requests, provide individuals with access to their data, rectify inaccuracies, and fulfill requests for erasure in accordance with privacy laws.

Furthermore, privacy laws mandate implementing robust security measures to protect personal data. Data inventories help organizations identify their personal data types and all associated security requirements. By clearly understanding the data they process, organizations can implement appropriate technical and organizational safeguards to shield personal data from unauthorized access, loss, or alteration, thus ensuring compliance with security obligations under privacy laws.

13.1.7 Data Destruction (Lesson Video)

Transcript:

When you dispose of a computer, selling used hardware, or erasing important information, it's crucial to destroy all of the data on a device. It's not enough to delete the data. Reformatting the hard drive is not even sufficient. If other people can access your computer, they can use data remanence, the residual representation of erased data, to recover information. You must damage the hardware so badly that the remanence is gone.

And don't paper documents. Even in our digital world, a lot of sensitive information is available in print, and you must destroy that information, too. If someone obtains your data, it might be used to crack passwords, sell identification information, or steal trade secrets. That's why it's essential to be familiar with the tools, techniques, and government laws that apply to data destruction.
There are wrong ways to destroy data and then there is the right way. You must use specialized tools to destroy data remanence, so information is impossible to retrieve. Data destruction tools use one of three techniques.
Software programs write random code over and over on top of the original data so that anything that can be recovered from the hardware is just gibberish.
Hardware tools damage hard disk drives so severely that it is impossible to retrieve even pieces of data from the drive.
There are also techniques that destroy the ink on paper so that it's impossible to read.
Let's look at sanitization programs first. Sanitization programs wipe data off your computer without damaging hardware. This is commonly called purging, wiping, or clearing. These techniques remove data from the hard disk and then write over it so many times that the data remanence is destroyed.
Media sanitization options do not destroy the hard disk for future use. There are many data sanitization software programs you can use to purge information, and some of them are even free. Each has pros and cons; carefully consider the media type you are purging, the level of data sensitivity, your budget, and other factors when you select a sanitization software.
The following options severely damage hardware storage so that the device can never be used again.
One technique is running a hard disk through a disk shredder, physically destroying the drive. Disk shredders are effective, but also relatively expensive.
Some organizations without a disk shredder will destroy a disk by hammering or drilling holes into the disk platters. Drilling and hammering are relatively inexpensive and fairly efficient. However, they are not as effective as a shredder.
Another option is hiring a company to destroy disks. Many companies that provide this service will also provide a certificate of destruction.
Pulverizing is like shredding, except that it uses a punch press or hammer system to crush a hard disk into a pile of metal confetti. Many companies that provide disk shredding also provide pulverizing.
One really effective technique is degaussing. Degaussing purges the hard disk by exposing it to an incredibly high magnetic pulse that destroys all the data on the disk. It also ruins the motors inside the drive.
Remember how we mentioned that you also need to destroy paper documents that contain sensitive information? Most offices contain paper shredders, which shred documents into very small pieces, as small as one-seventeenth of an inch which are then

mixed with other trash. But these paper pieces still contain ink. There are two methods that destroy the ink on the paper: burning and pulping.
Burning documents is a pretty simple process. Build a small fire somewhere legal and safe, then use metal tongs to burn your documents one by one or a few at a time. It's important to ensure that each document is turned into ash, if sensitive information escapes the flames and flies away, it might fall into the wrong hands.
Burning blank papers with your critical documents makes things more confusing if someone tries to piece them back together. Be sure to sift through all the remains and make sure every bit of paper has turned to ash before you dispose of the burnt remnants.
Pulping is a way of removing all traces of ink from paper by using chemicals and then mashing the paper into pulp. Since these chemicals can ruin carpet and clothing, you should perform this process outside and use protective gloves.
First, you need to remove the documents from envelopes, unfold them, and place them in a large trash can filled with nine parts water and one part liquid bleach. Second, completely submerge and soak the documents in the solution for 24 hours. Next, use a broomstick or a heavy-duty mixer, such as a paint turbine mixer, to mash the paper into a pliable liquid pulp.
After that, spread the pulp out to dry in the sun. Lastly, place the dried pulp into garbage bags and throw it in the trash.
Now that we have discussed ways to permanently dispose of information, let's switching gears to talk about the legal requirements of data destruction, starting with data sovereignty. Data sovereignty is the concept that information stored in binary digital form is subject to government laws and regulations. Data sovereignty laws are created to maintain privacy and prevent foreign countries from subpoenaing, or searching, another's data. Data sovereignty laws evolve rapidly as cloud services and other new storage options emerge.
It's imperative that you stay current with data sovereignty laws because every business must comply with state and federal data destruction regulations. Data sovereignty laws protect citizens from identity theft, companies from security breaches, and clients from privacy issues. They specify when, how, and what data you are allowed to destroy. Requirements vary depending on location, so make sure you comply with the most up-to-date version of both state and federal regulations, and remember, they frequently change to keep up with the pace of technology. No matter where you live in the US, you will have to follow laws dictated by three government acts: HIPAA, FACTA, and FISMA.
HIPAA stands for Health Insurance Portability and Accountability Act. HIPAA protects medical records and personal health information. Companies that provide healthcare insurance handle HIPAA-protected information. And, of course, companies that provide health-related services also handle HIPAA-protected information.
FACTA, the Fair and Accurate Credit Transactions Act, was created to protect against identity theft. It applies to the disposal of consumer reports and related information. FACTA includes credit reports, credit scores, employment history information, check writing history, insurance claims, residential or tenant history, and medical history. Every business handles FACTA-protected information, and every business must comply with FACTA laws.
FISMA, the Federal Information Security Management Act, protects government information. It is primarily concerned with proper data destruction and has detailed disposal requirements.
Regardless of which destruction method you use and what type of information you destroy, you need to carefully document each disposal process with a certificate of destruction, or COD. A COD is usually necessary to satisfy internal security reporting requirements as well as legal reporting requirements.
A COD should specify the device and the method and date of destruction. This document should also provide a chain of custody (each person who had possession of the data and the length of time they possessed the data).
If you destroy a storage device yourself, you should create your own COD. You can find COD templates on the internet. If you contact a service provider to destroy your devices, request that they provide a COD for each device they process. Most organizations, especially military and government agencies, retain CODs as a permanent record of each device's destruction.
Companies should have data retention policies and procedures for storing and destroying information. By law, organizations must retain certain information for specified time periods, and your business probably has additional records to keep and maintain. Always make absolutely certain that it is okay to destroy papers, data, and drives.
Remember, proper data destruction is essential to securing information. Not only must you use specialized tools and techniques to damage the data beyond recovery, you must also comply with all necessary laws and company policies, including creating and maintaining a certificate of destruction. Following these guidelines will protect your organization from many security and legal problems.

13.1.8 Data Destruction Facts

When a device is no longer needed, it often contains residual data, potentially sensitive information, and system configurations that could be exploited. A thorough and systematic decommissioning process ensures that all data is securely erased or overwritten to reduce the risk of exposure.

This lesson covers the following topics:

Decommissioning

Secure data destruction

Asset disposal

Decommissioning

Decommissioning processes play a vital role in supporting security within an organization. Decommissioning also involves resetting devices to their factory settings and eliminating any residual settings. Updating inventory records during decommissioning is also important to maintain an accurate account of active assets and support compliance requirements that mandate accurate asset tracking and secure disposal.

Decommissioning hardware securely is essential as hardware often stores sensitive data on internal drives and retains potentially exploitable configuration and user data. Data sanitization is a critically important step in the decommissioning process. It describes how all data on the device or removable media is securely erased to ensure no recoverable data remains.

Additionally, the device should be reset to its factory settings via a management console or other utility, eliminating any residual system configurations or settings that could pose a security risk. Disposing of physical equipment often warrants the physical destruction of some internal components like memory modules, hard disk drives (HDD), solid-state disks (SSD), M.2, or other storage modules, especially if they have stored sensitive data. In some scenarios, a professional disposal service specializing in certified secure disposal of electronic components may be the most appropriate choice.

The final step in the decommissioning process involves documentation and updating inventory records to reflect that the device has been decommissioned. This step ensures an accurate asset record and compliance with security standards and regulations.

For example, a multifunction network printer's decommissioning steps include sanitization of stored print jobs, scanned documents, and (potentially) fax transmissions; wiping stored network credentials and configuration data; performing a full factory reset; secure disposal or destruction of physical components; and asset inventory updates to reflect the device's status.

Secure Data Destruction

Several common circumstances may necessitate data destruction within an organization to ensure security, compliance, and proper management of resources. At the end of a data retention period, organizations must destroy data in accordance with internal policies and external regulations while optimizing storage resources. Legal and regulatory compliance, such as adhering to the General Data Protection Regulation (GDPR) or the Health Insurance Portability and Accountability Act (HIPAA), also requires the deletion or destruction of specific data when it is no longer needed or if requested by the data subject. Periodically destroying obsolete or outdated data can help maintain efficient storage utilization and reduce the risk of data breaches.

Additionally, when decommissioning storage devices or systems, destroy any stored data before disposal or repurposing to prevent unauthorized access to sensitive information.

Asset Disposal

Asset disposal/decommissioning concepts focus on the secure and compliant handling of data and storage devices at the end of their lifecycle or when they are no longer needed. Some important concepts include the following:

Concept	Description
Sanitization	Refers to the process of removing sensitive information from storage media to prevent unauthorized access or data breaches. This process uses specialized techniques, such as data wiping, degaussing,

Concept	Description
	or encryption, to ensure the data becomes irretrievable. Sanitization is particularly important when repurposing or donating storage devices, as it helps protect the organization's sensitive information and maintains compliance with data protection regulations.
Destruction	Involves the physical or electronic elimination of information stored on media, rendering it inaccessible and irrecoverable. Physical destruction methods include shredding, crushing, or incinerating storage devices, while electronic destruction involves overwriting data multiple times or using degaussing techniques to eliminate magnetic fields on storage media. Destruction is a crucial step in the decommissioning process and ensures that sensitive data cannot be retrieved or misused after the disposal of storage devices.
Certification	Refers to the documentation and verification of the data sanitization or destruction process. This often involves obtaining a certificate of destruction or sanitization from a reputable third-party provider, attesting that the data has been securely removed or destroyed in accordance with industry standards and regulations. Certification helps organizations maintain compliance with data protection requirements, provides evidence of due diligence, and reduces the risk of legal liabilities. Certifying data destruction without third-party involvement can be challenging, as the latter provides an impartial evaluation.

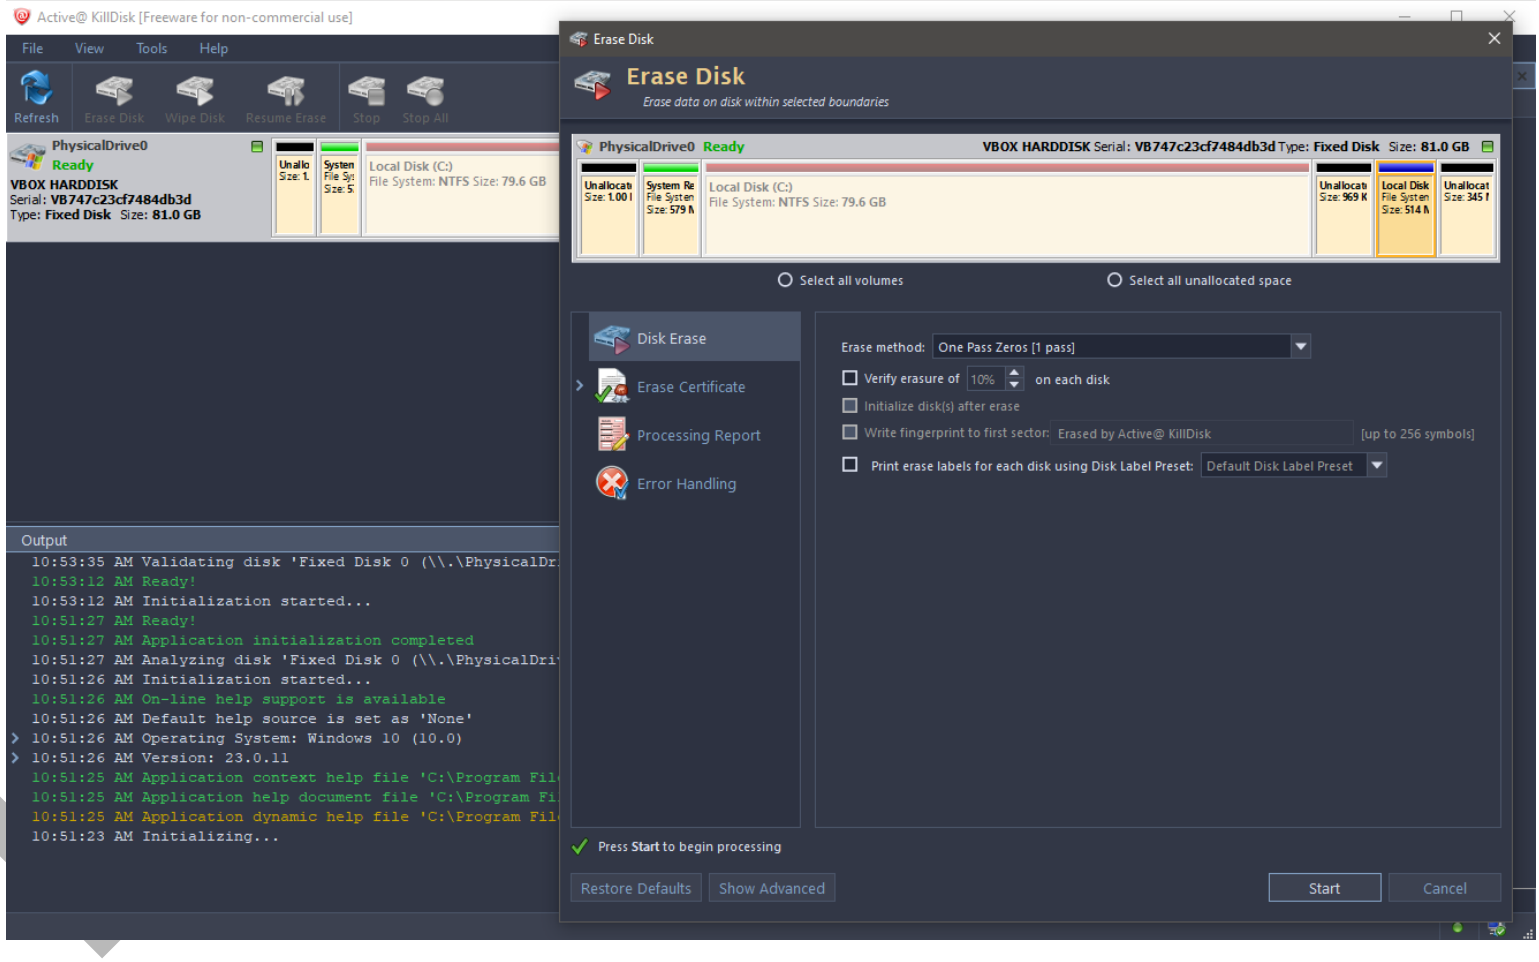

Active KillDisk data wiping software. (Screenshot used with permission from LSoft Technologies, Inc.)

Files deleted from a magnetic-type hard disk are not fully erased. Instead, the sectors containing the data are marked as available for writing, and the data they contain are only removed as new files are added. Similarly, the standard Windows format tool will only remove references to files and mark all sectors as usable. For this reason, the standard method of sanitizing an HDD is called **overwriting** . This can be performed using the drive's firmware tools or a utility program. The most basic type of overwriting is called zero filling, which sets each bit to zero. Single pass zero filling can leave patterns that can be read with specialist tools. A more secure method is to overwrite the content with one pass of all zeros, then a pass of all ones, and then a third pass in a pseudorandom pattern. Some federal agencies require more than three passes. Overwriting can take considerable time, depending on the number of passes required.

13.1.9 Privacy and Data Sensitivity Concepts

13.1.10 Practice Questions (Section Quiz)

q_breaches_data_breach_secp8

Which of the following can be consequences of a data breach? (Select three.)

Answers:

***Reputation damage**

***Fines**

***Identity theft**

Escalation

Financial penalties

Legal liabilities

Reputational damange

Explanation:

Data breaches can reflect negatively on a business and its ability to protect sensitive consumer information. Breaches can make consumers less likely to use the business again, damaging the company's reputation.

Fines are money that a court of law or another authority has levied on a company or a person as a penalty for wrongdoing. Fines are a consequence of data breaches.

Identity theft is intentionally using someone else's identity to gain a financial advantage or benefit without the other person's knowledge. Identity theft can be a consequence caused by breaches of usernames, passwords, or personally identifiable information (PII).

The following are not normally consequences of a data breach:

Escalation is the process of moving a problem or concern to a higher level of authority or deeper expertise.

Noncompliance can result in legal sanctions such as financial penalties, legal liabilities, reputational damage, and loss of customer trust.

q_breaches_escalation_secp8

As the chief information security officer (CISO) of a multinational corporation, you discover a significant data breach that has potentially exposed sensitive proprietary information.

What is the MOST crucial action you should take to mitigate the consequences of this breach?

Answers:

Notify the customers about the breach.

Pay the fines associated with the breach.

***Escalate the issue internally and externally to bring in experts for investigation.**

Work on repairing the company's reputation damage.

Explanation:

Escalating the issue internally and externally is the correct answer. In the case of a breach involving sensitive proprietary information, it is crucial to bring in experts for investigation as quickly as possible. This can help determine the extent of the breach, identify the perpetrators, and take necessary steps to prevent further damage.

Notifying customers about the breach is an important step when customer data is compromised. However, in this scenario, the breach involves proprietary information, not customer data. Therefore, the immediate action should be to escalate the issue to bring in experts for investigation.

Paying the fines associated with the breach is a consequence of a data breach, not an immediate action to mitigate its effects. Fines are typically levied after an investigation and determination of the company's compliance with data protection laws.

Repairing the company's reputation damage is a long-term process that comes after addressing the immediate needs of the breach, such as escalating the issue and investigating the breach. While it's an important aspect of recovery, it's not the most immediate action to take following a breach.

q_breaches_fines_01_secp8

A tech startup has just suffered a data breach where sensitive customer financial data leaked.

The chief executive officer (CEO) has an immediate concern about the tangible penalty the company will face due to violating data protection regulations.

What is the CEO primarily concerned with in this situation?

Answers:

Privacy policy updates

Reputational damage

***Fines**

Security infrastructure overhaul

Explanation:

Regulatory authorities impose fines as immediate, tangible penalties on companies violating specific regulations.

Privacy policy updates encompass changes in the company's stated data handling practices. Despite their importance, they do not serve as an immediate penalty for violating data protection regulations.

Non-compliance results in significant but intangible reputational damage, causing harm to the company's image among customers and stakeholders. However, this concern is not the CEO's immediate focus.

To address a data breach, a security infrastructure overhaul involves completely revamping the company's security systems. However, this is not an immediate, tangible penalty.

q_breaches_fines_02_secp8

A board of directors receives a memorandum that two departments in the organization violate federal regulations.

What could the organization receive that would monetarily impact them if sanctioned?

Answers:

***Fines**

Reputational damage

Loss of license

Indemnification

Explanation:

Sanctions refer to penalties, disciplinary actions, or measures imposed due to noncompliance with laws, regulations, or rules. Regulators may impose substantial fines, which can amount to millions or even billions of dollars, depending on the severity of the violation.

Not directly impacted by sanctions, noncompliance can harm an organization's reputation, eroding customer trust, decreasing business opportunities, and potentially losing contracts or partnerships.

Not impacted by sanctions, noncompliance with software usage can result in contractual noncompliance, leaving work products vulnerable.

Indemnification is not directly associated with the impacts of a sanction but may result in the noncompliant part assuming liability for damages caused by a security breach or data loss.

q_breaches_indemnification_secp8

A cyber team evaluates areas that pose more risk of becoming noncompliant.

What is the ramification of indemnification?

Answers:

***Liability for damages caused by data loss**

Modifying code or distributing software

Exceeding permitted installations

Unauthorized sharing or usage

Explanation:

Not associated with software licensing but applicable to this scenario, indemnification occurs as the result of damages caused by a security breach or data loss.

Noncompliance with software licensing requirements can result in the revocation of usage rights and other consequences, such as fines. Violations of license agreements, such as modifying code or distributing software without authorization, constitute noncompliance with the owner of the software.

Other forms of noncompliance for software licensing include the organization exceeding the permitted number of installations.

The unauthorized sharing or other unauthorized usage can also jeopardize a loss of software licensing due to noncompliance.

q_breaches_iptheft_secp8

Your organization has discovered that an overseas company has reverse-engineered and copied your main product and is now selling a counterfeit version.

Which of the following BEST describes the type of consequence your organization has suffered?

Answers:

Fines

Escalation

Reputation damage

***IP theft**

Explanation:

Intellectual property (IP) is the lifeblood of companies. When their IP is stolen, they lose competitive advantage. The internet has made the world smaller, and companies are now competing with others from around the globe. When IP is stolen through a data breach, it is often sold to competing companies. Many of the companies are overseas and not subject to laws in the United States, making them difficult to prosecute. This allows reverse-engineering or direct copying of IP and gives thieves an undeserved revenue source. This also floods the market with counterfeit goods.

A company's reputation determines if people invest, if consumers buy a product or service, or if foreign governments even allow a certain company to do business in their jurisdiction. This scenario does not describe this.

Fines can be levied against an organization as a result of a data breach. This scenario does not describe this.

Escalation can be separated into two categories, which are internal escalation and external escalation. Internal escalation is part of a company's incident-response plan. External escalation is when experts need to be brought in from the outside to investigate, provide legal counsel, or even enforce laws.

q_breaches_noncompliance_secp8

What describes the impacts associated with contractual noncompliance?

Answers:

***It can result in a breach or termination of an agreement or indemnification.**

It can grant certain individuals to challenge credit data on their personal reports.

It can include financial penalties, legal liabilities, and loss of customer trust.

Numerous governing bodies, such as regulatory authorities, can oversee the organization.

Explanation:

Contractual noncompliance occurs when organizations fail to meet agreed requirements, possibly resulting in a breach of contract, termination of the contract, or indemnification.

The "right to be forgotten" is a fundamental principle outlined in the General Data Protection Regulation (GDPR) that grants individuals the right to request the erasure or deletion of their personal data under certain circumstances.

Common ramifications for noncompliance include legal sanctions such as financial penalties, legal liabilities, and loss of customer trust. Sanctions refer to penalties, disciplinary actions, or measures imposed due to noncompliance with laws, regulations, or rules.

Sanctions enforcement is the responsibility of governing bodies, regulatory authorities, or organizations overseeing the specific domain in which the noncompliance occurred.

q_breaches_notification_secp8

As the chief information security officer (CISO) of a large corporation, you discover a significant data breach that has potentially exposed sensitive customer data.

What is the most immediate action you should take to mitigate the consequences of this breach?

Answers:

Escalate the issue internally and externally to bring in experts for investigation.

***Notify the customers about the breach.**

Pay the fines associated with the breach.

Work on repairing the company's reputation damage.

Explanation:

Notifying the customers about the breach is the correct answer. Each of the 50 states now have laws on breach notifications. Most of the laws have common elements. For example, a company must immediately disclose a breach to its customers in writing. This allows customers to take immediate action to protect themselves, which can help mitigate the overall impact of the breach.

Escalating the issue internally and externally is an important step in managing a data breach. However, it is not the most immediate action. The first step should be to notify the affected parties so they can take protective measures.

Paying the fines associated with the breach is a consequence of a data breach, not an immediate action to mitigate its effects. Fines are typically levied after an investigation and determination of the company's compliance with data protection laws.

Repairing the company's reputation damage is a long-term process that comes after addressing the immediate needs of the breach, such as notifying customers and investigating the breach. While it's an important aspect of recovery, it's not the most immediate action to take following a breach.

q_breaches_reputation_secp8

Your organization has suffered a data breach, and it was made public. As a result, stock prices have fallen, as consumers no longer trust the organization.

Which of the following BEST describes the type of consequence your organization has suffered due to the breach?

Answers:

Identity theft

Notifications

***Reputation damage**

IP theft

Explanation:

This scenario best describes an organization's reputation damage from a data breach. A company's reputation determines if people invest, if consumers buy a product or service, or if foreign governments even allow a certain company to do business in their jurisdiction. A company lives or dies by its revenues and investments, so a breach that exposes client data directly affects the way consumers and investors spend their money. A data breach can cause stock prices to fall, and falling stock prices lead to selloffs and permanent damage.

Notifications are usually sent out following a data breach. This scenario does not describe this.

Intellectual property (IP) is the lifeblood of companies. When their IP is stolen, they lose competitive advantage. This scenario does not describe this.

Identity theft is when a breach occurs and personal information is stolen. The affected individual or entity is forced to do hours of work to correct someone else's mistake.

q_breaches_sanctions_secp8

A board of directors convenes to discuss reports that the tech department was not meeting legal regulations.

What are the impacts associated with sanctions? (Select two.)

Answers:

***It can include financial penalties, legal liabilities, and loss of customer trust.**

***It can be overseen by numerous governing bodies, such as regulatory authorities.**

It can result in a breach or termination of an agreement, or indemnification.

It can grant certain individuals with the ability to challenge credit data on their personal reports.

It can result in information being read, modified, or deleted without authorization.

Explanation:

Typical ramifications for noncompliance include legal sanctions such as financial penalties, legal liabilities, and loss of customer trust. Sanctions refer to penalties, disciplinary actions, or measures imposed due to noncompliance with laws, regulations, or rules.

Sanctions enforcement is the responsibility of governing bodies, regulatory authorities, or organizations overseeing the specific domain in which the noncompliance occurred.

Contractual noncompliance occurs when organizations fail to meet agreed requirements, possibly resulting in a breach of contract, termination of the contract, or indemnification.

The "right to be forgotten" is a fundamental principle outlined in the General Data Protection Regulation (GDPR) that grants individuals the right to request the erasure or deletion of their personal data under certain circumstances.

A data breach occurs when information is read, modified, or deleted without authorization. It is not an impact of sanctions.

q_breaches_software_secp8

A cyber technician purchases new software to monitor employee computer usage after the company announces a new work-from-home program.

Which risks could lead to noncompliance with software agreements? (Select three.)

Answers:

***Modifying code or distributing software**

***Exceeding permitted installations**

***Unauthorized sharing or usage**

Liability for damages caused by data loss

Liability caused by a security breach

Data transferred to a network or storage media

Loss or disclosure of personal and sensitive data

Explanation:

Noncompliance with software licensing requirements can result in the revocation of usage rights and other consequences, such as fines. Violations of license agreements, such as modifying code or distributing software without authorization, constitute noncompliance with the owner of the software.

Other forms of noncompliance for software licensing include the organization exceeding the permitted number of installations.

The unauthorized sharing or other unauthorized usage can also jeopardize a loss of software licensing due to noncompliance.

Not associated with software licensing includes:

Indemnification occurs as the result of damages caused by a security breach or data loss.

Data transferred to a network or storage media

Loss or disclosure of personal and sensitive data

q_info_class_compliance_secp8

A multinational corporation handles regulated data. What are the key considerations for handling and protecting the data to ensure compliance?

Answers:

***Data subject to legal and regulatory requirements: privacy, security, compliance**

Data requiring encryption: secure storage, transmission, adherence to standards

Financial and trade secret data: reporting, risk management, non-disclosure agreements

Data requiring user awareness training: content filtering, web security, authentication

Explanation:

Regulated data refers to information governed by laws and regulations. Key considerations include privacy, security, and compliance with relevant requirements.

While encryption may be a consideration, it does not fully define regulated data, and the key considerations extend beyond encryption.

Regulated data encompasses more than financial and trade secret data; the key considerations mentioned are not comprehensive. Financial and trade secret data would not be key considerations for the organization to ensure compliance.

While user awareness training may be relevant, it does not fully capture the concept of regulated data and the broader key considerations for compliance.

q_info_class_confidential_secp8

A financial organization is currently handling a document that contains sensitive customer information, including financial details and social security numbers.

According to data classifications, how should the financial organization categorize this data?

Answers:

***Confidential data**

Trade secret data

Proprietary data

Restricted data

Explanation:

The information is highly sensitive, and only approved persons within the organization and possibly trusted third parties under the non-disclosure agreement (NDA) should view it. Since this data includes sensitive customer information, including financial details and social security numbers, the organization should classify it as confidential.

Trade secret data is a type of information that holds significant commercial value and is kept secret for competitive advantage.

Proprietary information or intellectual property is information the company creates and owns, typically about the products or services they make or perform.

While this data is sensitive, the term "restricted" usually applies to information that could cause significant harm to individuals, organizations, or national security if disclosed or accessed by unauthorized individuals.

q_info_class_financial_secp8

A multinational company is preparing to submit its annual statements and needs to share confidential reports with its global offices. The data includes the company's revenues, profits, and other sensitive information.

What data category do these annual statements fall under?

Answers:

***Financial data**

Trade secret data

Human-readable data

Regulated data

Explanation:

The annual financial statement, including all related financial information, falls into this category. Legal and financial data is sensitive and confidential and can significantly impact an organization's reputation, legal standing, and financial stability.

While the annual financial statement might contain sensitive and valuable information, it does not contain trade secrets such as proprietary formulas, processes, or methods that would give the company a competitive advantage.

Although the annual financial statement could classify as human-readable, as it can be easily understood and interpreted by humans, this category does not reflect the financial and legal significance of the data.

Regulated data typically includes personally identifiable information (PII), credit card details, and health records directly regulated by laws and regulations.

q_info_class_human_readable_01_secp8

A company's marketing team creates an infographic for an upcoming campaign, presenting detailed statistics and forecasts to better inform their audience about the product's performance.

In terms of data types, how would the marketing team classify this infographic?

Answers:

***Human-readable data**

Trade secret data

Legal and financial data

Regulated data

Explanation:

The infographic's design is to be easily understood and interpreted by humans without the need for additional processing or translation, making it human-readable data.

While the data contained in the infographic could be proprietary or confidential, it is for public marketing purposes, so it does not qualify as trade secret data.

Although the infographic might contain some performance statistics, the primary purpose of the infographic is not for legal compliance or financial reporting. Therefore, it does not fall under the legal and financial data type.

Regulated data typically includes personally identifiable information (PII) or other data types directly regulated by specific laws and regulations. Since the infographic does not contain such information, it is not regulated data.

q_info_class_human_readable_02_secp8

What type of data is information that can easily be understood and interpreted without additional processing or translation?

Answers:

Regulated data

Trade secrets

***Human-readable data**

Non-human-readable data

Explanation:

Human-readable data is information in a format directly comprehensible to humans without additional processing or translation.

Regulated data refers to categories of information subject to legal or regulatory requirements for their handling, storage, and protection. It does not necessarily mean humans can easily understand or interpret it without additional processing or translation.

Trade secrets refer to confidential business information that provides a competitive edge. Like regulated data, the trade secret category does not indicate the data is easily understandable or interpretable by humans.

Humans cannot easily understand or interpret non-human-readable data in its raw form, and it often requires additional processing or translation.

q_info_class_legal_secp8

An organization is preparing to file for a patent for its innovative product design. They have gathered all necessary information, including detailed descriptions of the design, how it differs from existing designs, and why it is eligible for patent protection.

Under which data type would this information fall?

Answers:

***Legal and financial data**

Human-readable data

Trade secret data

Regulated data

Explanation:

The patent application, including all related designs and justifications, falls into this category. Legal data include legal agreements, court records, intellectual property filings, and other legal documents that are of significant importance to an organization's legal standing.

While the patent information may be human-readable, this classification does not fully encapsulate the specific legal nature of the information.

Although the information about the design could be a trade secret until the organization files the patent, during the patent process, the information will be publicly accessible and not trade secret data.

This category typically includes personally identifiable information (PII) and other data types directly regulated by specific laws and regulations, and a patent does not fall into this category.

q_info_class_nist_secp8

Companies can categorize information security and cybersecurity tasks into five distinct functions.

Which regulatory concept or entity offers guidance and frameworks for this classification?

Answers:

***National Institute of Standards and Technology (NIST)**

General Data Protection Regulation (GDPR)

Payment Card Industry Data Security Standard (PCI DSS)

Center for Internet Security (CIS)

Explanation:

Companies can organize their information security and cybersecurity tasks into the following: identify, protect, detect, respond, and recover. This follows the framework of the National Institute of Standards and Technology.

The General Data Protection Regulation stipulates companies cannot collect, process, or retain personal data without the individual's informed consent.

The Payment Card Industry Data Security Standard defines the safe handling and storage of financial information.

The Center for Internet Securityis a not-for-profit organization founded partly by The SANS Institute. It publishes the 20 CIS Controls.

q_info_class_pii_data_secp8

You are a data analyst at a financial institution. During a routine audit, you discover that a colleague has been storing customers' personally identifiable information (PII) in a non-encrypted, local database for ease of access.

The colleague insists that this practice is safe as the database is password-protected and only accessible from their work computer.

What should you do?

Answers:

Ignore the situation since the database is password-protected and only accessible from your colleague's work computer.

***Report the situation to your supervisor and recommend that the PII data be moved to an encrypted, secure server.**

Confront your colleague and insist that they delete the database immediately.

Access the database yourself to verify the extent of the PII data stored.

Explanation:

Reporting the situation to your supervisor ensures that the issue is handled appropriately and that the PII data is moved to a secure location is the most appropriate action. It also ensures that the company is in compliance with data protection laws and regulations.

Password protection and limited access do not provide sufficient security for PII data. There are still risks such as potential hacking or physical theft of the work computer.

While it's important to address the issue with your colleague, insisting that they delete the database immediately could result in loss of important data. It's better to report the situation to your supervisor who can decide on the best course of action.

Accessing the database yourself could potentially violate privacy policies and regulations. It's best to report the situation to your supervisor and let them handle it.

q_info_class_pii_secp8

If you lose your wallet or purse and it ends up in the wrong hands, several pieces of information could be used to do personal harm to you. These pieces of information include the following:

Name and address

Driver license number

Credit card numbers

Date of birth

Which of the following classifications does this information fall into?

Answers:

***Personally identifiable information**

Proprietary information

Private internal information

Private restricted information

Explanation:

Personally identifiable information (PII) is information that can be used on its own or with other information to identify, contact, or locate a single person. This information includes:

Full name (if not common)

Home address

Email address (if private from an association/club membership, etc.)

National identification number

Passport number

IP address (when linked, but it is not PII by itself in US)

Vehicle registration plate number

Driver's license number

Face, fingerprints, or handwriting

Credit card numbers

Digital identity

Date of birth

Birthplace

Genetic information

Telephone number

Login name, screen name, nickname, or handle

Proprietary information is information that a company wishes to keep confidential. Private internal information is restricted to individuals within the organization and can include personnel records, financial records, and customer lists. Private restricted information is restricted to limited authorized personnel within the organization and can include trade secrets, strategic information, and highly sensitive information.

q_info_class_proprietary_01_secp8

A software company designs a new feature for its product involving the creation and storage of new algorithms and methods that give the product a competitive advantage. The company wants to appropriately classify this information within its data management system.

What would be the MOST fitting classification for this data?

Answers:

Public

Confidential

Critical

***Proprietary**

Explanation:

Proprietary information or intellectual property (IP) refers to information the company creates and owns, typically concerning the products or services it makes or performs.

Public or unclassified data refers to information with no restrictions on viewing and no risk to an organization if someone were to disclose it.

While this data is sensitive, the confidential classification typically applies to information meant for viewing only by approved persons within the organization or by trusted third parties under a non-disclosure agreement.

The term critical generally applies to data or systems whose loss would severely impact the operations or survival of an organization. Although the information is vital to the company, proprietary is more suitable for the information the company owns.

q_info_class_proprietary_02_secp8

A company is very protective of its intellectual material. The fear of a breach by a curious public or competitors is an ongoing concern.

As a result, the company put in place a dedicated server containing related highly sensitive data.

Applying knowledge of data types and labels, which data type is the company is protecting?

Answers:

***Proprietary**

Public

Private

Confidential

Explanation:

Proprietary information, or intellectual property (IP), is information a company creates, typically about the products or services it makes or performs.

Public information bears no classification and has no restrictions on viewing. Companies should be careful not to release information to the public without ensuring the risk of that data.

Classified material (i.e., private, restricted, internal use only, official use only) restricts viewing to the owner organization or third parties under a non-disclosure agreement (NDA).

Confidential (or low) information is a military classification scheme that is highly sensitive and intended for viewing only by approved persons within the organization and possibly by trusted third parties under NDA. However, it may not necessarily be confidential or top secret.

q_info_class_public_01_secp8

A governmental organization is preparing to release census data that includes the total population, age distribution, and employment rates.

According to the data classifications, how should the organization categorize this data?

Answers:

***Public data**

Confidential data

Proprietary data

Restricted data

Explanation:

Public data refers to information that has no restrictions on viewing. The governmental organization's census data, which presents no risk to the organization if disclosed, fits this definition.

Confidential data is highly sensitive and meant for viewing only by approved persons within the organization and possibly by trusted third parties under a non-disclosure agreement (NDA).

Proprietary information or intellectual property is information the company creates and owns, typically about the products or services that it makes or performs.

Restricted data refers to sensitive information that requires stringent controls and limited access due to its highly confidential nature. Restricted data does not fit the description of the census data intended to be publicly released.

q_info_class_public_02_secp8

A national park posts information about its flora and fauna on its website. This information does not contain any personally identifiable information or sensitive government data.

How should the park service classify this data?

Answers:

***Public**

Regulated

Confidential

Private

Explanation:

Public data can include any data that is free to access, such as information about the park's flora and fauna. Likewise, it does not contain any personally identifiable information or sensitive government data.

Regulated data does not encapsulate public data, such as nonsensitive information intended for the public like the park's data.

Confidential data is sensitive information. Data about the park's flora and fauna does not need this type of protection, making this classification inappropriate.

Private classification would require the information about the park's flora and fauna not to be freely available to the public.

q_info_class_readable_secp8

A multinational corporation handles human-readable and non-human-readable data.

What are the implications for security operations and controls?

Answers:

***Security measures for human-readable data: monitoring, user awareness, DLP, and content filtering**

Security measures for human-readable data: monitoring, user awareness, encryption, and secure data exchange

Security measures for non-human-readable data: monitoring, access controls, intrusion detection/prevention, and code/application security

Security measures for non-human-readable data: encryption, access controls, intrusion detection/prevention, and secure data exchange

Explanation:

These measures are directly applicable to protect data formats easily understood and interpreted by humans, such as documents and web pages. Monitoring helps detect and respond to potential threats, while user awareness enhances security practices and data loss prevention (DLP) and content filtering mitigate risks associated with sensitive data.

Encryption and secure data exchange are more relevant to non-human-readable data formats, not human-readable data.

Monitoring is more applicable to human-readable data formats, not non-human-readable data.

Encryption and secure data exchange are relevant to non-human-readable data formats, but access controls and intrusion detection/prevention are also applicable.

q_info_class_regulated_01_secp8

An organization prepares to store and handle a data type that includes sensitive personal information, such as healthcare records and social security numbers. This data is subject to specific laws and regulations concerning its protection and use.

What category does this data type fall under?

Answers:

Trade secrets

Human-readable data

***Regulated data**

Legal and financial data

Explanation:

Regulated data refers to specific categories of information subject to legal or regulatory requirements regarding their handling, storage, and protection, which typically includes sensitive or personally identifiable information (PII), such as healthcare records and social security numbers.

Trade secret data refers to valuable, confidential information that gives a business a competitive advantage.

Human-readable data is information that humans can easily understand and interpret without additional processing or translation.

While this data type is highly sensitive and confidential, it pertains to documents, contracts, financial statements, balance sheets, audit reports, tax records, financial transactions, and other such legal and financial documents.

q_info_class_regulated_02_secp8

The IT department of a healthcare provider maintains a database containing personal health information for its patients.

Which classification BEST suits this type of data?

Answers:

***Regulated**

Critical

Public

Nonsensitive

Explanation:

The classification is that of regulated data since personal health information (PHI) is subject to strict legal and compliance regulations to protect patient privacy, such as Health Insurance Portability and Accountability Act (HIPAA) in the United States.

Critical data could impact the organization's operations or ability to function effectively. While PHI is important, it is not necessarily critical to the survival of the healthcare organization.

Public data is information that the general public can freely disseminate without restrictions. PHI is sensitive and confidential; thus, it cannot be public.

PHI is sensitive and confidential, not public (nonsensitive). Public data is information that is NOT confidential, private, sensitive, regulated, or has any other restrictions.

q_info_class_report_document_secp8

As a cybersecurity analyst at a government agency, you come across a document labeled as "Top Secret Compartmentalized".

You have a Top Secret clearance, but you're not sure if you're authorized to view this document. The document could potentially contain information that would help you resolve a current cybersecurity issue.

What should you do?

Answers:

Open the document. You have Top Secret clearance, so you should be able to view all Top Secret documents.

***Do not open the document. Instead, report the document to your supervisor and ask for guidance.**

Open the document, but only read the sections that you believe are relevant to your current cybersecurity issue.

Ignore the document. If it was meant for you, it would have been sent to you directly.

Explanation:

Report the document to your supervisor to ask for guidance is the most appropriate action to take. You should not access the document without confirming that you are authorized to do so. Reporting the document to your supervisor ensures that you are following protocol and not potentially breaching any security policies.

Even though you have Top Secret clearance, the "Compartmentalized" label means the document has very specific and limited distribution. You should not assume you have access to all Top Secret documents.

Even if you believe the document could be useful, you should not access it without the necessary authorization. Doing so could result in a security breach and potential disciplinary action.

Ignoring the document does not address the potential security issue. It's important to report the document to your supervisor to ensure proper handling.

q_info_class_trade_secret_01_secp8

A tech startup develops a unique algorithm that provides a significant competitive edge in the market. To maintain this edge, the startup needs to ensure the highest level of protection for this information.

How should this startup categorize and handle this unique algorithm?

Answers:

***The startup should categorize the algorithm as a trade secret and protect it using non-disclosure agreements (NDAs).**

The startup should categorize the algorithm as regulated data and follow compliance regulations.

The startup should categorize the algorithm as human-readable data and make it accessible to all employees.

The startup should categorize the algorithm as legal and financial data and keep it secure for financial reporting.

Explanation:

A trade secret is confidential information that gives a business a competitive advantage. The unique algorithm developed by the startup falls into this category. Using non-disclosure agreements can help safeguard the confidentiality of this trade secret.

Regulated data refers to information subject to legal or regulatory requirements regarding handling, storage, and protection, such as personally identifiable information (PII) or health records.

While the algorithm might be human-readable, making it accessible to all employees is not ideal as it increases the risk of leakage or theft.

Legal and financial data relates to critical data for legal compliance, financial reporting, and risk management. While the algorithm might be critical to the startup, it is not considered legal or financial data.

q_info_class_trade_secret_02_secp8

A software development company has just completed an innovative, proprietary software that it plans to release on the market soon. This software includes unique code and intellectual property that they developed in-house and believe is a crucial asset.

What data category does the software's unique code fall under?

Answers:

***Trade secret data**

Regulated data

Legal and financial data

Human-readable data

Explanation:

Trade secrets refer to valuable, confidential information that gives a business a competitive advantage. The unique, proprietary software code developed by the company would fall into this category.

Regulated data usually refers to personal data or data categories directly regulated by laws and regulations, such as financial information or health records. The unique code developed by a software company would not typically fall into this category.

While the company's financials reflect the value of the unique software code, the code itself would not classify as financial or legal data.

Even though software code could technically be human-readable, classifying it as such would not capture the strategic value and need for confidentiality associated with the proprietary software code.

q_info_class_types_secp8

The government and military use the following information classification system:

Unclassified

Sensitive But Unclassified

Confidential

Secret

Top Secret

Drag each classification on the left to the appropriate description on the right.

Answers:

Unclassified

Sensitive But Unclassified

Confidential

Secret

Top Secret

Explanation:

The government and military use the following information classification system:

Unclassified: information that can be accessed by the public and poses no security threat.

Sensitive But Unclassified: information that, if disclosed, could cause some harm but not a national disaster.

Confidential: information that is the lowest level of classified information used by the military. It allows restriction of release under the Freedom of Information Act. Release of this information could cause damage to military efforts.

Secret: information that, if disclosed, could cause severe and permanent damage to military actions.

Top Secret: information that is the highest level of classified information used by the military. If Top Secret information is released, it poses a grave threat to national security.

q_priv_data_resp_controller_secp8

A tech startup is rolling out a new platform that collects user feedback to improve its software. The company manages the data collection, decides how to process it, and has subcontracted a third-party analytics firm to analyze the feedback.

Which role is the startup playing?

Answers:

***Data controller**

Data subject

Data processor

Data custodian

Explanation:

The controller determines the purposes and means of processing personal data. In this case, the startup decides how the data processes and manages its collection, making it the controller.

The data subject is the individual who provides personal data to an organization. In this scenario, the users giving feedback are the data subjects.

The processor processes personal data on behalf of the controller. The third-party analytics firm, which analyzed the feedback, plays this role.

The data custodian is responsible for maintaining and protecting the data. While the startup does handle the data, its primary role is not merely safeguarding it but deciding how to use and process it.

q_priv_data_resp_custodian_secp8

A small department at a company manages a server, separate from IT, for data access and backup purposes.

What role does the department fulfill?

Answers:

***Data custodian**

Data owner

Data controller

Data processor

Explanation:

The data custodian role manages the system on which the data assets reside. This role includes enforcing access control, encryption, and backup/recovery measures.

The data owner is usually in a senior (executive) role with ultimate responsibility for maintaining the information asset's confidentiality, integrity, and availability.

A data controller is responsible for determining why and how employees can collect, keep, and use data and ensuring that these purposes and means are lawful.

A data processor is an entity the data controller engages to assist with technical collection, storage, or analysis tasks.

q_priv_data_resp_data_retention_secp8

A company tasks a data team with decommissioning an organizational data lake.

What are the tenets associated with the data retention concept? (Select the best three options.)

Answers:

***It is often based on legal, regulatory, or operational requirements.**

***It ensures that organizations maintain compliance with relevant regulations and minimize the risk of breaches.**

***It refers to policies and practices governing the storage and preservation of information within the organization for a set period.**

It refers to documenting and verifying the data sanitization or destruction process.

It determines the purposes and means of processing personal data.

It processes personal data on behalf of the data controller.

It provides the right to request access to personal data and obtain information about how it is being processed.

Explanation:

Data retention refers to the policies and practices governing the storage and preservation of information within an organization for a specified period.

In addition, data retention policies are often based on legal, regulatory, or operational requirements and dictate when and how data should be securely deleted or destroyed.

Proper data retention practices ensure that organizations maintain compliance with relevant regulations, optimize storage resources, and minimize the risk of data breaches or unauthorized access to sensitive information during the disposal and decommissioning process.

The following are not tenants associated with the data retention concept:

Certification refers to documenting and verifying the data sanitization or destruction process. It is not part of the data retention concept.

A data controller is the entity or organization that determines the purposes and means of processing personal data.

The data processor processes personal data on behalf of the data controller.

Data subjects hold certain rights. One of these is the right of access, meaning that data subjects have the right to request access to their personal data and obtain information about how it is being processed.

q_priv_data_resp_data_roles_secp8

A healthcare organization is strengthening its data protection framework to ensure compliance with local and international regulations. One focus area is clearly defining the roles and responsibilities between the data controllers and processors, as this impacts the overall management and protection of sensitive data.

In this scenario, which two statements accurately outline the responsibilities of the data controller and the data processor regarding data protection? (Select two.)

Answers:

***Data controller--determines the purposes for which data is processed**

***Data processor--processes data on behalf of the controller**

Data controller--performs day-to-day operations on data

Data processor--decides the purpose of data processing

Both roles have the same responsibilities.

Explanation:

A data controller is primarily responsible for determining the purpose of processing the data, which guides the organization.

On the other hand, the data processor processes the data on behalf of the data controller, following the guidelines the controller provides.

The responsibility of performing day-to-day operations, such as data storage or management, typically falls under the role of the data processor, not the controller.

The data processor does not decide the purpose of data processing. This responsibility belongs to the data controller. The processor operates under the direction of the controller.

The roles in this (and most scenarios) are different, and not the same.

q_priv_data_resp_gdpr_01_secp8

Which of the following accurately reflects the responsibilities of a data processor under data protection laws such as the General Data Protection Regulation (GDPR)?

Answers:

The data processor determines the purposes and means of processing personal data.

The data processor has independent decision-making power over personal data.

The data processor is directly responsible for obtaining consent from data subjects.

***The data processor processes personal data on behalf of the data controller.**

Explanation:

The data processor processes personal data on behalf of the data controller and acts under the authority and instructions of the data controller. The data processor is not allowed to make decisions alone regarding the processing of the data.

The responsibility of determining the purposes and means of processing personal data lies with the data controller, not the data processor.

The data processor does not have independent decision-making power over personal data. The data processor can only process data when the data controller instructs the data processor to do so.

The responsibility of obtaining consent from data subjects lies with the data controller, not the data processor.

q_priv_data_resp_gdpr_02_secp8

Under the General Data Protection Regulation (GDPR), how soon must an organization report a breach of personal data?

Answers:

***Within 72 hours of becoming aware of the breach**

Within 24 hours of becoming aware of the breach

Within 96 hours of becoming aware of the breach

Within 48 hours of becoming aware of the breach

Explanation:

The General Data Protection Regulation (GDPR) requires that an organization must report a breach of personal data to the relevant supervisory authority within 72 hours of discovering it unless the breach is unlikely to risk the rights and freedoms of natural persons.

GDPR requires that organizations must report personal data breaches within 72 hours, not the shorter time frames of 24 or 48 hours.

A time frame of 96 hours is longer than the period specified by GDPR. The regulation stipulates that organizations must report breaches within 72 hours of becoming aware of them. Delaying notification beyond the mandated timeframe could result in penalties.

q_priv_data_resp_gdpr_03_secp8

An organization evaluates the legal implications of failing to protect privacy data after experiencing a breach.

What level of influence does the GDPR have regarding legal implications?

Answers:

***Global**

Regional

National

Local

Explanation:

The General Data Protection Regulation (GDPR) in the European Union has had a substantial impact globally by setting high privacy and data protection standards.

At the regional level, privacy data protections consist of local environments offering regional guidance on data privacy concerns.

At the national level, data protection authorities enforce privacy laws, oversee compliance, and have the authority to investigate data breaches, taking legal action against organizations that fail to protect privacy data or violate individuals' privacy rights.

At the local level, data privacy protection gets limited to smaller areas, enforced only within a confined jurisdiction, and not at regional, national, or global levels.

q_priv_data_resp_gdpr_ccpa_secp8

A multinational firm headquartered in San Francisco, California, serves customers from various countries, including European Union countries.

The company collects, processes, and stores substantial amounts of personal data.

With which of the following legal regulations must the company's governance committee ensure compliance?

Answers:

General Data Protection Regulation (GDPR) only

California Consumer Privacy Act (CCPA) only

***Both General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA)**

Neither General Data Protection Regulation (GDPR) nor California Consumer Privacy Act (CCPA)

Explanation:

GDPR requires companies to protect the personal data and privacy of EU citizens for transactions that occur within the EU. At the same time, CCPA provides California residents with specific rights regarding their personal information.

Although the company serves customers from the European Union and thus needs to comply with the GDPR, it also has its headquarters in California and is subject to the CCPA.

Because the company also serves customers in the European Union, it needs to ensure compliance with the GDPR.

As the company headquarters is in California and serves customers from the European Union, it must comply with both the GDPR and the CCPA; thus, this option is incorrect.

q_priv_data_resp_inventories_01_secp8

When evaluating privacy laws, what provides a comprehensive overview of the types of handled data?

Answers:

***Data inventories**

Data retention

Attestation and acknowledgment

Due diligence

Explanation:

Data inventories provide a comprehensive overview of the types of handled data, the purposes for processing, the legal basis, and the recipients of the data to ensure transparency and accountability.

Data retention is another area impacted by privacy laws. Organizations must retain personal data only for as long as necessary to fulfill the intended purpose or as required by law.

Attestation and acknowledgment require individuals or entities to formally acknowledge their understanding of compliance obligations and commitment to adhere to them through signed agreements, policy acknowledgments, and training activities.

Due diligence in data protection describes the comprehensive assessment and evaluation of an organization's data protection practices and measures.

q_priv_data_resp_inventories_02_secp8

How do data inventories assist organizations in maintaining records of collected data?

Answers:

***It provides a comprehensive overview of the types of handled data.**

It is an established timeline that requires organizations to keep documentation.

It requires individuals or entities to announce their understanding of compliance obligations formally.

It is the comprehensive assessment and evaluation of an organization's data protection practices.

Explanation:

Data inventories provide a comprehensive overview of the types of handled data, the purposes for processing, the legal basis, and the recipients of the data to ensure transparency and accountability.

Data retention is another area impacted by privacy laws. Organizations must retain personal data only for as long as necessary to fulfill the intended purpose or as required by law.

Attestation and acknowledgment require individuals or entities to formally acknowledge their understanding of compliance obligations and commitment to adhere to them through signed agreements, policy acknowledgments, and training activities.

Due diligence in data protection describes the comprehensive assessment and evaluation of an organization's data protection practices and measures.

q_priv_data_resp_owner_secp8

You are the data privacy officer (DPO) at a large corporation.

A new project involving the processing of personal data is about to start. The project manager approaches you for guidance on roles and responsibilities related to data privacy.

Which of the following advice would be MOST appropriate to give?

Answers:

The project manager should act as the data owner, controller, and processor.

***The project manager should act as the data owner, and appoint separate individuals as the data controller and processor.**

The project manager should act as the data controller, and appoint separate individuals as the data owner and processor.

The project manager should act as the data processor, and appoint separate individuals as the data owner and controller.

Explanation:

The project manager should act as the data owner, and appoint separate individuals as the data controller and processor is the correct answer. The data owner is usually a senior stakeholder responsible for a dataset. The data controller and processor should be separate individuals to ensure checks and balances. The controller decides the purpose and procedure for processing personal data, while the processor carries out the daily processing of personal data.

The project manager should act as the data owner, controller, and processor is not correct. It is not advisable for one person to take on all these roles as it could lead to conflicts of interest and potential breaches in data privacy.

The project manager should act as the data controller, and appoint separate individuals as the data owner and processor is not correct. The project manager, who is overseeing the project, would be more suited to the role of data owner. The roles of data controller and processor should be assigned to individuals who are involved in the daily handling and processing of data.

The project manager should act as the data processor, and appoint separate individuals as the data owner and controller is not correct. The project manager, given their seniority and oversight role, would be more suited to the role of data owner. The roles of data controller and processor should be assigned to individuals who are involved in the daily handling and processing of data.

q_priv_data_resp_right_forgotten_secp8

A healthcare provider must maintain comprehensive patient records while ensuring the privacy of individuals' information.

How can the provider navigate legal requirements for data retention with respect to patients who request that their information is removable?

Answers:

Establishing a specific data retention policy

Maintaining extended data inventory

***Complying with the right to be forgotten**

Regularly reviewing and updating privacy policies

Explanation:

Involving procedures that enable patients to request the deletion of their data, complying with the right to be forgotten requires the healthcare provider to evaluate whether these requests align with the current legal framework.

Separate from individuals' rights to control their personal information, establishing a specific data retention policy presents its own different considerations.

While maintaining an extended data inventory can aid in complying with various regulations, it does not address how organizations should handle patient requests for data removal.

Regularly reviewing and updating privacy policies focuses on ensuring legal compliance rather than serving as a solution to conflicts between legal retention obligations and individual data removal requests.

q_priv_data_resp_roles_secp8

Security governance relies heavily on specially designed and interdependent roles. Each role has unique responsibilities that contribute to effective security oversight and control.

What are some of these roles? (Select three.)

Answers:

***Owner**

***Controller**

***Processor**

Maintenance custodian

Administrator

Engineer

Technician

Explanation:

A high-ranking employee (i.e., a director or vice president) typically holds the owner role and is ultimately responsible for ensuring data is appropriately protected.

The controller role closely relates to the General Data Protection Regulation (GDPR) and identifies the purposes, conditions, and means of processing personal data.

The processor is responsible for processing personal data on behalf of the controller and often represents Cloud Service Providers (CSP), although they could also be vendors and business partners.

The following are NOT security governance roles:

A maintenance custodian performs the upkeep and care of a building and is responsible for its cleanup. On the other hand, a security governance custodian is responsible for the safe custody, transport, and storage of the data. They are two entirely different roles.

Administrator, engineer, and technician are networking roles.

q_priv_data_scenario_secp8

You are a data protection officer at a tech company. The company is planning to launch a new product that collects user data for personalized recommendations. The data collected includes user names, addresses, contact information, and their browsing history on the company's platform.

Which of the following statements BEST describes how the company should handle this data?

Answers:

The company should treat all collected data as confidential data since it is collected for business purposes.

The company should treat all collected data as privacy data since it can be used to identify specific individuals.

***The company should treat user names, addresses, and contact information as privacy data, and browsing history as confidential data.**

The company should not differentiate between privacy data and confidential data and should apply the same protection measures to all collected data.

Explanation:

User names, addresses, and contact information are privacy data as they can be used to identify specific individuals. Browsing history on the company's platform is considered confidential data as it pertains to the company's proprietary information is the correct answer.

While the data is collected for business purposes, it includes personally identifiable information (PII) such as user names, addresses, and contact information, which should be treated as privacy data.

While user names, addresses, and contact information are privacy data, browsing history on the company's platform is considered confidential data as it pertains to the company's proprietary information.

Privacy data and confidential data have different legal and ethical considerations. Privacy data pertains to personal information and individual privacy rights, while confidential data is concerned with safeguarding information from unauthorized access, use, or disclosure to maintain business competitiveness and protect intellectual property. Therefore, different protection measures should be applied to each type of data.

q_data_destroy_facts_certificate_01_secp8

A healthcare organization is preparing to decommission several servers containing sensitive patient information. The organization wants to ensure that it securely disposes of the data on these servers and properly documents this process.

What should the organization primarily focus on to ensure secure data disposal and regulation compliance?

Answers:

***Obtain a certificate of destruction or sanitization from a third-party provider.**

Sell the servers as soon as possible to reclaim some of the initial investment.

Keep the servers in storage indefinitely as a backup in case of data loss.

Pass the servers to the IT team for random allocation among employees.

Explanation:

Obtaining certification proves that the organization securely eliminated data in compliance with industry standards and regulations. It also reduces the risk of legal liabilities and serves as evidence of due diligence.

Failing to ensure the secure destruction of the sensitive data before selling the servers could lead to serious data breaches and regulatory noncompliance.

Keeping the servers in storage indefinitely could increase the risk of a data breach over time and unnecessarily ties up storage space that could have better uses.

This action does not address the primary issue of securely destroying sensitive data. It also creates potential security risks, as employees could access or misuse sensitive data.

q_data_destroy_facts_certificate_02_secp8

A security specialist is drafting a memorandum on secure data destruction for the organization after a recent breach.

What benefit does the certification concept offer when evaluating appropriate disposal/decommissioning?

Answers:

***It refers to the documentation and verification of the data sanitization or destruction process.**

It refers to policies and practices governing the storage and preservation of information within the organization for a set period of time.

It is often based on legal, regulatory, or operational requirements.

It ensures that organizations maintain compliance with relevant regulations and minimize breach risks.

Explanation:

Certification refers to documenting and verifying the data sanitization or destruction process.

Data retention, not certification, refers to the policies and practices governing the storage and preservation of information within an organization for a specified period.

Data retention policies are often based on legal, regulatory, or operational requirements and dictate when and how organizations should securely delete or destroy data.

Proper data retention practices ensure that organizations maintain compliance with relevant regulations, optimize storage resources, and minimize the risk of data breaches or unauthorized access to sensitive information during the disposal and decommissioning process.

q_data_destroy_facts_combo_secp8

You are a cybersecurity expert at a large corporation. Your company has just decommissioned a data center and you are tasked with ensuring the secure destruction of data on thousands of hard drives. The data includes highly sensitive information.

Which of the following solutions would be BEST for destroying the data?

Answers:

Overwrite the data on the drives and then reuse them within the company.

Use a third-party service to degauss the drives.

Physically destroy the drives by pulverizing them.

***Use a combination of overwriting and degaussing before disposing of the drives.**

Explanation:

Using a combination of overwriting and degaussing before disposing of the drives is the best solution in this scenario. This option provides a high level of security by first overwriting the data (making it unreadable) and then degaussing the drives (rendering them unusable). This two-step process reduces the chance of data recovery to near zero. While the drives cannot be reused, the level of data sensitivity and the scale of the task may warrant this approach.

While overwriting the data can make it unreadable and allows for the medium to be reused, it may not be the most secure method for highly sensitive data. There's a small chance that some data could potentially be recovered.

Degaussing changes the magnetic field of the drive, rendering it unusable and the data unrecoverable. This is a secure method of data destruction for highly sensitive data. However, it means the drives cannot be reused, which could be a waste of resources given the large number of drives.

While physically destroy the drives by pulverizing them would ensure the data is unrecoverable, it is not the most environmentally friendly option. It also doesn't allow for the possibility of reusing or recycling the drives.

q_data_destroy_facts_decomission_secp8

In preparation for taking over inventory for their division, an engineer reviews the records for the previous inventory count. The engineer notes that a few items were on the previous inventory count that they could not find during their initial search.

What records should they review prior to looking for the missing items?

Answers:

***Decommissioning**

Access control

Configuration enforcement

Monitoring

Explanation:

The final step in the decommissioning process involves updating inventory records to reflect decommissioned devices. The decommissioning records explain why the property was on the previous inventory.

Access control relates to permissions granted to individuals, software, networks, or systems to only allow the completion of authorized actions by these entries. It is not a part of the inventory process.

Configuration enforcement ensures that systems and devices within a network have the same configuration, allowing for easier identification of compromised systems. It is not a part of the inventory process.

Monitoring plays a vital role in the hardening process. Additionally, it can provide valuable data for compliance and auditing purposes. It is not a part of the inventory process.

q_data_destroy_facts_degaus_01_secp8

Which of the following is the LEAST reliable means of cleaning or purging media?

Answers:

OS low-level formatting

***Degaussing**

Overwriting every sector with alternating 1s and 0s

Drive controller hardware-level formatting

Explanation:

The least reliable means to clean or purge media is degaussing. Degaussing is the use of strong magnetic fields to remove stored information from a drive. Unfortunately, user error and equipment failure often results in only partially cleaned media.

Various forms of formatting (such as OS low-level formatting and drive controller hardware-level formatting) are not perfect, but they are often more reliable than degaussing. Overwriting every sector with alternating 1s and 0s can be effective if performed multiple times (such as 60 or more).

q_data_destroy_facts_degaus_02_secp8

You are the head of the IT department at a large corporation. Your company has just completed a major data migration project and you now have a significant number of old hard drives that contain sensitive company data. You need to ensure this data is completely destroyed to prevent any potential data breaches.

Which method would you choose and why?

Answers:

Overwrite the data on the drives and then reuse them within the company.

***Use a third-party service to degauss the drives.**

Physically destroy the drives by pulverizing them.

Overwrite the data on the drives and then sell them to a recycling company.

Explanation:

Use a third-party service to degauss the drives is the correct answer. Degaussing changes the magnetic field of the drive, rendering it unusable and the data unrecoverable. This is a secure method of data destruction for highly sensitive data. However, it means the drives cannot be reused.

While overwriting the data can make it unreadable and allows for the medium to be reused, it may not be the most secure method for highly sensitive data. There's a small chance that some data could potentially be recovered.

While physically destroying the drives would ensure the data is unrecoverable, it is not the most environmentally friendly option. It also doesn't allow for the possibility of reusing or recycling the drives.

Overwriting the data on the drives and then selling them to a recycling company has the same potential security issue as option A. Additionally, selling the drives to a third party, even a recycling company, could potentially expose the company to a data breach if the overwriting process was not 100% effective.

q_data_destroy_facts_degaus_03_secp8

A financial services company is decommissioning many servers that contain highly sensitive financial information. The company's data protection policy stipulates the need to use the MOST secure data destruction methods and comply with strict regulatory requirements.

The company also has a significant environmental sustainability commitment and seeks to minimize waste wherever possible.

What should the company's primary course of action be during this process?

Answers:

***Degaussing the servers, rendering the data irretrievable, followed by reselling or recycling the servers after certification**

Incinerating the servers, as it's the most effective method of data destruction

Overwriting the data on the servers multiple times, then disposing of the servers without any certification

Storing the servers indefinitely in a secure location to avoid any risk of data leakage

Explanation:

This method would ensure the secure destruction of sensitive data and compliance with regulatory requirements. Also, the company can uphold its environmental sustainability commitment by reselling or recycling.

Incinerating the servers is an effective data destruction method, but it would not align with the company's commitment to environmental sustainability. Additionally, it may not be the most cost-effective solution, especially if the servers are repurposeable or recyclable.

Overwriting the data on the servers multiple times, then disposing of the servers without certification would not meet strict regulatory requirements.

Storing the servers indefinitely in a secure location would be inefficient and unnecessary. This method would occupy physical space and resources indefinitely and pose a constant risk of potential data breaches.

q_data_destroy_facts_destruction_01_secp8

When cleaning out the server closet, a company discovers a box of old disk drives.

When considering which disposal method to use, what are the characteristics associated with the destruction concept? (Select two.)

Answers:

***It involves the physical or electronic elimination of information stored on media, rendering it inaccessible and irrecoverable.**

***Its methods include shredding, crushing, or incinerating storage devices.**

It refers to removing sensitive information from storage media to prevent unauthorized access or data breaches.

Its process uses specialized techniques, such as data wiping, degaussing, or encryption.

It refers to copying files to other media to keep in a secure safe.

Explanation:

Destruction involves the physical or electronic elimination of information stored in media, rendering it inaccessible and irrecoverable.

Destruction methods include shredding, crushing, or incinerating storage devices, while electronic destruction involves overwriting data multiple times or using degaussing techniques to eliminate magnetic fields on storage media.

Sanitization refers to removing sensitive information from storage media to prevent unauthorized access or data breaches.

Sanitization uses specialized techniques, such as data wiping, degaussing, or encryption, to ensure the data becomes irretrievable. Sanitization is particularly important when repurposing or donating storage devices.

Copying the files to other media to keep in a secure safe is not a key concept of destruction.

q_data_destroy_facts_destruction_02_secp8

A healthcare organization is retiring an old database server that housed sensitive patient information. It aims to ensure that this information is completely irretrievable.

What key process should the organization prioritize before disposing of this server?

Answers:

***Secure destruction of all data stored on the server**

Certification of the server's functionality

Preservation of all data for future reference

Repurposing of the server without any modifications

Explanation:

Secure destruction involves the physical or electronic elimination of information stored on media, rendering it inaccessible and irrecoverable. It is a crucial step before disposing of any storage devices containing sensitive information.

Certification is important, but it usually refers to documenting and verifying the data sanitization or destruction process. In this context, it is not the initial step before disposal.

Preserving sensitive data when a storage device decommissions increases the risk of data breaches and non-compliance with data protection regulations.

Repurposing storage devices without first sanitizing them could lead to unauthorized access to sensitive information. In this context, the organization must destroy the data before the disposal or repurposing of a server.

q_data_destroy_facts_destruction_03_secp8

An organization has decommissioned several laptops used for handling sensitive data.

Which of the following should be the primary step to ensure data security and compliance with regulations before repurposing or disposing of these devices?

Answers:

Conducting a hardware audit

Removing all software applications

Deleting all user accounts

***Initiating a secure data destruction process**

Explanation:

Initiating a secure data destruction process ensures the irretrievable deletion of sensitive data on the laptops' storage media, preventing unauthorized access or data breaches.

Conducting a hardware audit is an important part of asset management and could be useful in tracking the decommissioned laptops. However, it does not directly ensure the security of the sensitive data previously stored on the devices.

Removing all software applications from the laptops could make some data less accessible, but it will not ensure the secure deletion of sensitive data.

Deleting all user accounts can remove some user-specific data and settings from the laptops but will not securely delete or sanitize all the sensitive data stored on the devices.

q_data_destroy_facts_pulv_01_secp8

When you dispose of a computer or sell used hardware, it is crucial that none of the data on the hard disks can be recovered.

Which of the following actions can you take to ensure that no data is recoverable?

Answers:

***Damage the hard disks so badly that all data remanence is gone.**

Reformat all the hard disks in the computer.

Delete all files from all the hard disks in the computer.

Encrypt all data on the hard disks.

Explanation:

When you dispose of a computer, sell used hardware, or erase important information, it's crucial to destroy all of the data on a device. It's not enough to delete the data. Reformatting the hard drive is also not sufficient. If other people can access the computer, they can use data remanence (the residual representation of erased data) to recover information. You must damage the hardware so badly that the remanence is gone.

q_data_destroy_facts_pulv_02_secp8

Which of the following data destruction techniques uses a punch press or hammer system to crush a hard disk?

Answers:

***Pulverizing**

Shredding

Pulping

Purging

Degaussing

Explanation:

The following are various ways to destroy data:

Pulverizing: like shredding except that it uses a punch press or hammer system to crush a hard disk into a pile of metal confetti.

Burning: the method of building a small fire somewhere legal and safe. Use metal tongs to burn your documents one by one or a few at a time. It's important to ensure that each document is turned into ash. If sensitive information escapes the flames and flies away, it might fall into the wrong hands.

Shredding: running a hard disk through a disk shredder, physically destroying the drive.

Pulping: a way of removing all traces of ink from paper by using chemicals and then mashing the paper into pulp. Since these chemicals can ruin carpet and clothing, you should perform this process outside and use protective gloves.

Degaussing: purges the hard disk by exposing it to a high magnetic pulse that destroys all of the data on the disk. This also ruins the motors inside the drive.

Purging: the removal of sensitive data, ensuring that it cannot be reconstructed by any known technique.

Wiping: a software-based method of overwriting data to completely destroy all electronic data residing on a hard disk drive or other digital media. Wiping uses 0s and 1s to overwrite data onto all sectors of the device. The data is rendered unrecoverable and achieves data sanitization.

q_data_destroy_facts_sanitize_01_secp8

A tech company is in the process of decommissioning a fleet of old servers. It wants to ensure that sensitive data stored on these servers is fully eliminated and are not accesible in the event of unauthorized attempts.

What primary process should the company implement before disposing or repurposing these servers?

Answers:

***Sanitizing the servers**

Deleting all the files on the servers

Moving the servers to a secure storage location

Selling the servers immediately

Explanation:

Sanitizing the servers uses specialized techniques, such as data wiping, degaussing, or encryption, to ensure the data becomes irretrievable.

Deleting all files on the servers is ineffective because deleting files from a magnetic-type hard disk does not fully erase them.

Moving the servers to a secure location and keeping them untouched does not address the need to remove sensitive data from the servers. If someone accesses the servers without authorization, it would increase the data breach risk.

Selling the servers immediately without considering the stored data is risky as it could lead to unauthorized access to sensitive data, legal liabilities, and reputational damage.

q_data_destroy_facts_sanitize_02_secp8

A financial institution is preparing to decommission a number of its old servers. The servers contain sensitive customer data that needs proper handling to prevent unauthorized access or data breaches.

Which strategy should the institution primarily employ to ensure the data on these servers stays irretrievable?

Answers:

***Carry out a sanitization process that includes multiple passes of overwriting and degaussing.**

Leave the data on the servers, as the system will eventually overwrite it.

Use a basic method of overwriting, such as zero filling, once.

Physically destroying the servers is necessary.

Explanation:

The sanitization process ensures that the data on the servers is irretrievable, protecting sensitive customer data from unauthorized access or data breaches.

Leaving the data on the servers until it eventually overwrites would not be a secure option. This approach would leave the data vulnerable to unauthorized access or data breaches.

Single-pass zero filling can leave patterns readable with specialist tools, so multiple passes become necessary for proper sanitization.

Physically destroying the servers without considering any data wiping process could leave data accessible. Physical destruction should be a final step after the company thoroughly sanitizes the data electronically.

q_data_destroy_facts_sanitize_03_secp8

Upon receiving new storage media drives for the department, an organization asks a software engineer to dispose of the old drives.

When considering the various methods, what processes does sanitization involve? (Select two.)

Answers:

***It refers to the process of removing sensitive information from storage media to prevent unauthorized access or data breaches.**

***Its process uses specialized techniques, such as data wiping, degaussing, or encryption.**

It involves the physical or electronic elimination of information stored in media, rendering it inaccessible and irrecoverable.

Its methods include shredding, crushing, or incinerating storage devices.

It involves the documentation and verification of the destruction process.

Explanation:

Sanitization refers to removing sensitive information from storage media to prevent unauthorized access or data breaches.

Sanitization uses specialized techniques, such as data wiping, degaussing, or encryption, to ensure the data becomes irretrievable. Sanitization is particularly important when repurposing or donating storage devices.

Destruction involves the physical or electronic elimination of information stored in media, rendering it inaccessible and irrecoverable.

Destruction methods include shredding, crushing, or incinerating storage devices, while electronic destruction involves overwriting data multiple times or using degaussing techniques to eliminate magnetic fields on storage media.

Certification refers to the documentation and verification of the data sanitization or destruction process.

13.2 Personnel Policies

As you study this section, answer the following questions:

How can an onboarding process improve the security of an organization?

Why should employees be required to sign employment agreements?

Why is it important to conduct an exit interview?

What security issues must be identified and addressed during the onboarding phase of a third-party relationship?

What is the role of the Service Level Agreement (SLA)?

Why are policies important for organizational security?

The key terms for this section include:

Term	Definition
Onboarding	The process of bringing in a new employee, contractor, or supplier.
Offboarding	The process of ensuring that all HR and other requirements are covered when an employee leaves an organization.
Acceptable use policy (AUP)	A policy that governs employees' use of company equipment and Internet services. ISPs may also apply AUPs to their customers.

Code of conduct	Professional behavior depends on basic ethical standards, such as honesty and fairness. Some professions may have developed codes of ethics to cover difficult situations; some businesses may also have a code of ethics to communicate the values it expects its employees to practice.
Clean desk policy	An organizational policy that mandates employee work areas be free from potentially sensitive information; sensitive documents must not be left out where unauthorized personnel might see them.
Computer-based training (CBT)	Training and education programs delivered using computer devices and e-learning instructional models and design.
Anomalous behavior recognition	Systems that automatically detect users, hosts, and services that deviate from what is expected, or systems and training that encourage reporting of this by employees.
Database encryption	Applying encryption at the table, field, or record level via a database management system rather than via the file system.
Data sovereignty	In data protection, the principle that countries and states may impose individual requirements on data collected or stored within their jurisdiction.

This section helps you prepare for the following certification exam objectives:

Exam	Objective
CompTIA Security+ SY0-701	2.5 Explain the purpose of mitigation techniques used to secure the enterprise. Least privilege 3.3 Compare and contrast concepts and strategies to protect data. Data types Regulated Intellectual property Legal information Financial information Data classifications General data considerations

Data states

Data at rest

Data in transit

Data in use

Data sovereignty

Geolocation

Methods to secure data

Encryption

4.1 Given a scenario, apply common security techniques to computing resources.

Mobile solutions

Bring your own device (BYOD)

4.2 Explain the security implications of proper hardware, software, and data asset management.

Disposal/decommissioning

Data retention 4.4 Explain security alerting and monitoring concepts and tools.

Tools

Data loss prevention (DLP)

4.6 Given a scenario, implement and maintain identity and access management.

Access controls

Least privilege

4.8 Explain appropriate incident response activities.

Training

Digital forensics

Legal hold

5.1 Summarize elements of effective security governance.

	Policies Acceptable use policy (AUP) Procedures Onboarding/offboarding 5.4 Summarize elements of effective security compliance. Privacy Legal implications Data inventory and retention 5.6 Given a scenario, implement security awareness practices. Phishing Recognizing a phishing attempt Anomalous behavior recognition Risky Unexpected Unintentional User guidance and training Policy/handbooks Situational awareness Removable media and cables Social engineering Reporting and monitoring Initial Recurring Development
TestOut Security Pro	5.0 Audit and Security Assessment

13.2.1 Personnel Policies (Lesson Video)

Transcript:

In this lesson, we're going to discuss personnel policies. Employees are a critical part of any successful business, but they also represent one of the greatest risks to your organization's security. Using proper employment practices helps mitigate this risk. They should encompass all aspects of employment, from pre-employment screenings to offboarding.

Employee security should start even before an employee starts work. To ensure that a potential employee is a low security risk, an organization should perform a background check. This check may include checking into a person's criminal history, credit history, and employment history. It could also involve drug testing, identity verification, or a review of their social media presence.
Once you've decided to hire an individual, you'll need to set up a work environment for them. This process is called onboarding. It primarily involves a lot of HR paperwork, but there are a few onboarding procedures that involve security professionals.
When you're setting up the new employee's user accounts, security credentials, and building access, make sure to assign the appropriate access rights. To do this, you'll need to know the new employee's job title, duties, and place in the organizational hierarchy.
It's also important that new employees sign necessary forms and documents. A non-disclosure agreement stipulates that employees can't reveal confidential or proprietary information to any non-employee. A non-compete agreement prevents former employees from working for competitors. And an acceptable use policy specifies how employees are allowed to use the information organization's network resources.
As we've discussed, employees are a serious security risk. Sometimes they're intentionally trying to cause harm, but more often, they're simply uninformed. One way to help secure your data and your network is to hold regular security trainings for your employees. These trainings should include topics like how to recognize security breaches and exploits when they occur, what to do in the event of a security breach, and sample scenarios with recommended responses.
Training should be ongoing and train employees on new security issues and refresh them on older topics. Ideally, you should use a combination of teaching techniques so you have a better chance of engaging everyone and helping them retain information.
In addition to classroom training, security methods can be taught through games, simulations, role-based hands-on training, or computer-based training.
Onboarding and training provide a great starting point, but you still need ongoing policies to ensure that employee performance and duties are secure.
So, how can we better secure our daily operations? First, let's look at separation of duties. Separation of duties restricts the amount of access or influence individuals control, which removes single points of failure within the organization.
For example, let's say an employee in accounting approves business expenses, verifies expenditures, signs checks, and withdraws funds. The risk of this employee embezzling company money is significantly multiplied because this one person performs all of these duties. In this situation, this person could easily become a single point of failure.
With separation of duties, expenditures and withdrawals should be verified by at least two sets of people. And maybe the CFO should sign the checks instead of the accountant. This way, we can eliminate this single point of failure.
You can also enhance security by implementing duty rotations. For example, after a couple of months, each person's responsibility shifts to another person, like a chore chart.
So, using our previous example, instead of signing checks, the CFO verifies and approves expenditures, while others handle withdrawals and signing checks. Duty rotations also heighten security by continually monitoring accountability during the rotation process.
Now let's move on to the principle of least privilege. The idea behind this principle is that employees only have access to information they absolutely need for their job function. Employees with too much access could either intentionally or unintentionally compromise security.
For example, administrative accounts should only be given to people who need them, and this access should only be used by employees who are actively performing administrative duties. A best practice is to log in with a standard user account and then only change to an administrative account to perform a task that requires it. When you're done with that task, you switch back to the standard user account.

It's never a good idea to be logged in as an administrator while browsing the web, reading emails, or using social media. An exploit coming through that internet connection could easily obtain elevated privileges on that workstation, which hugely increases the potential for damages.
Another practice to increase security is a mandatory vacation policy. The best way to implement mandatory vacations is to require employees to request vacation days in advance, like at the beginning of the year. The benefits of mandatory vacations are twofold. First, they allow employees to take a break from work, rest, and recharge. Second, they provide the security administrator the ability to detect fraudulent activities while employees are away. For example, if you notice that a reoccurring security breach stops while a particular employee is away on vacation, that could help you identify its cause.
The last step in employment is offboarding. Whether an employee resigns, retires, or is terminated, offboarding procedures keep the company safe as they depart.
First, you should disable or delete that employee's user account. An active, unused account allows easy access to your network. When you delete a user account, be sure not to delete that user's data in the process. Make sure their files are preserved. You want that data available for future use. So you may want to disable former employees' accounts rather than delete them.
You should also change security keys and require the return of access tokens, devices, and phones. And when an employee leaves the organization, Human Resources should always conduct an exit interview.
During this interview, it's important to address any concerns, dissatisfactions, or feedback from the employee. With this information, you'll have a better idea of why the employee is leaving and be able to address any remaining concerns. This is also a good time to remind the employee of any non-disclosure or non-compete agreements they may have signed during their onboarding process.
That's it for this lesson. In this lesson, we discussed the importance of personnel policies during the pre-employment process, onboarding, training, and daily operations, and we finished by discussing offboarding policies that protect the company and mitigate risks.

13.2.2 Personnel Policy Facts

Personnel policies play a vital role in establishing clear guidelines, expectations, and standards for employees. They provide a framework for effective management of human resources and maintain a fair, legally compliant, and productive work environment. Personnel policies promote consistency, clear communication, employee development, conflict resolution, and risk management. Organizations can enhance employee satisfaction, attract and retain talent, and mitigate legal and operational risks by implementing effective personnel policies.

This lesson covers the following topics:

- Personnel Management
- Conduct policies
- Training topics and techniques
- Reporting and monitoring

Personnel Management

Human Resources (HR) policies for personnel management are applied in three phases:

- **Recruitment (hiring)** — locating and selecting people to work in particular job roles. Security issues here include screening candidates and performing background checks.
- **Operation (working)** — it is often the HR department that manages the communication of policy and training to employees (though there may be a separate training and personal development department within larger organizations). As such, it is critical that HR managers devise training programs that communicate the importance of security to employees.

Termination or Separation (firing or retiring) — whether an employee leaves voluntarily or involuntarily, termination is a difficult process, with numerous security implications.

Background Checks

A background check determines that a person is who they say they are and are not concealing criminal activity, bankruptcy, or connections that would make them unsuitable or risky. Employees working in high confidentiality environments or with access to high-value transactions will obviously need to be subjected to a greater degree of scrutiny. For some jobs, especially federal jobs requiring a security clearance, background checks are mandatory. Some background checks are performed internally, whereas others are done by an external third party.

Onboarding

Onboarding at the HR level is the process of welcoming a new employee to the organization. The same sort of principle applies to taking on new suppliers or contractors. Some of the same checks and processes are used in creating customer and guest accounts.

As part of onboarding, the IT and HR function will combine to create an account for the user to access the computer system, assign the appropriate privileges, and ensure the account credentials are known only to the valid user. These functions must be integrated, to avoid creating accidental configuration vulnerabilities, such as IT creating an account for an employee who is never actually hired. Some of the other tasks and processes involved in onboarding include the following:

Secure Transmission of Credentials — creating and sending an initial password or issuing a smart card securely. The process needs protection against rogue administrative staff. Newly created accounts with simple or default passwords are an easily exploitable backdoor.

Asset Allocation — provision computers or mobile devices for the user or agree to the use of bring-your-own-device handsets.

Training/Policies — schedule appropriate security awareness and role-relevant training and certification.

Offboarding

An exit interview (or offboarding) is the process of ensuring that an employee leaves a company gracefully. Offboarding is also used when a project using contractors or third parties ends. In terms of security, there are several processes that must be completed:

Account Management — disable the user account and privileges. Ensure that any information assets created or managed by the employee but owned by the company are accessible (in terms of encryption keys or password-protected files).

Company Assets — retrieve mobile devices, keys, smart cards, USB media, and so on. The employee will need to confirm (and in some cases prove) that they have not retained copies of any information assets.

Personal Assets — wipe employee-owned devices of corporate data and applications. The employee may also be allowed to retain some information assets (such as personal emails or contact information), depending on the policies in force.

The departure of some types of employees should trigger additional processes to re-secure network systems. Examples include employees with detailed knowledge of security systems and procedures, and access to shared or generic account credentials. These credentials must be changed immediately.

Conduct Policies

Operational policies include privilege/credential management, data handling, and incident response. Other important security policies include those governing employee conduct and respect for privacy.

Acceptable Use Policy

Enforcing an acceptable use policy (AUP) is important to protect the organization from the security and legal implications of employees misusing its equipment. Typically, the policy will forbid the use of equipment to defraud, defame, or to obtain illegal material. It will prohibit the installation of unauthorized hardware or software and explicitly forbid actual or attempted snooping of confidential data that the employee is not authorized to access. Acceptable use guidelines must be reasonable and not interfere with employees' fundamental job duties or privacy rights. An organization's AUP may forbid use of Internet tools outside of work-related duties or restrict such use to break times.

Code of Conduct and Social Media Analysis

A code of conduct , or rules of behavior, sets out expected professional standards. For example, employees' use of social media and file sharing poses substantial risks to the organization, including threat of virus infection or systems intrusion, lost work time, copyright infringement, and defamation. Users should be aware that any data communications, such as email, made through an organization's computer system are likely stored within the system, on servers, backup devices, and so on. Such communications are also likely to be logged and monitored. Employers may also subject employees' personal social media accounts to analysis and monitoring, to check for policy infringements.

Rules of behavior are also important when considering employees with privileged access to computer systems. Technicians and managers should be bound by clauses that forbid them from misusing privileges to snoop on other employees or to disable a security mechanism.

Use of Personally Owned Devices in the Workplace

Portable devices, such as smartphones, USB sticks, media players, and so on, pose a considerable threat to data security, as they make file copy so easy. Camera and voice-recording functions are other obvious security issues. Network access control, endpoint management, and data loss prevention solutions can be of some use in preventing the attachment of such devices to corporate networks. Some companies may try to prevent staff from bringing such devices on-site . This is quite difficult to enforce, though.

Also important to consider is the unauthorized use of personal software by employees or employees using software or services that has not been sanctioned for a project (shadow IT). Personal software may include either locally installed software or hosted applications, such as personal email or instant messenger, and may leave the organization open to a variety of security vulnerabilities. Such programs may provide a route for data exfiltration, a transport mechanism for malware, or possibly software license violations for which the company might be held liable, just to name a few of the potential problems.

Clean Desk Policy

A clean desk policy means that each employee's work area should be free from any documents left there. The aim of the policy is to prevent sensitive information from being obtained by unauthorized staff or guests at the workplace.

Training Topics and Techniques

It is necessary to frame security training in language that end users will respond to. Education should focus on responsibilities and threats that are relevant to users. It is necessary to educate users about new or emerging threats (such as fileless malware, phishing scams, or zero-day exploits in software), but this needs to be stated in language that users understand. Using a diversity of training techniques helps to improve engagement and retention. Training methods include facilitated workshops and events,

one-on-one instruction and mentoring, plus resources such as computer-based or online training, videos, books, and blogs/newsletters.

User and Role-Based Training

Another essential component of a secure system is effective user training. Untrained users represent a serious vulnerability because they are susceptible to social engineering and malware attacks and may be careless when handling sensitive or confidential data.

Train users in secure behavior. (Image by dotshock © 123RF.com.)

Appropriate security awareness training needs to be delivered to employees at all levels, including end users, technical staff, and executives. Some of the general topics that need to be covered include the following:

- Overview of the organization's security policies and the penalties for noncompliance.
- Incident identification and reporting procedures.
- Site security procedures, restrictions, and advice, including safety drills, escorting guests, use of secure areas, and use of personal devices.
- Data handling, including document confidentiality, PII, backup, encryption, and so on.
- Password and account management plus security features of PCs and mobile devices.
- Awareness of social engineering and malware threats, including phishing, website exploits, and spam plus alerting methods for new threats.
- Secure use of software such as browsers and email clients plus appropriate use of Internet access, including social networking sites.

There should also be a system for identifying staff performing security-sensitive roles and grading the level of training and education required (between beginner, intermediate, and advanced, for instance). Note that in defining such training programs you need to focus on job roles, rather than job titles, as employees may perform different roles and have different security training, education, or awareness requirements in each role.

The NIST National Initiative for Cybersecurity Education framework (nist.gov/itl/applied-cybersecurity/nice) sets out knowledge, skills, and abilities (KSAs) for different cybersecurity roles. Security awareness programs are described in SP800-50 (nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-50.pdf).

Computer-Based Training and Gamification

Participants respond well to the competitive challenge of CTF events. This type of gamification can be used to boost security awareness for other roles too. Computer-based training (CBT) allows a student to acquire skills and experience by completing various types of practical activities:

Simulations — recreating system interfaces or using emulators so students can practice configuration tasks.

Branching scenarios — having students choose between options to find the best choices to solve a cybersecurity incident or configuration problem.

CBT might use video game elements to improve engagement. For example, students might win badges and level-up bonuses such as skills or digitized loot to improve their in-game avatar. Simulations might be presented so that the student chooses encounters from a map and engages with a simulation environment in a first person shooter type of 3D world.

Critical Elements for Security Awareness Training

Training employees about safe computer use is critical to protecting data and mitigating the risks associated with cyberattacks. (Image by rawpixel © 123RF.com.)

Phishing campaigns used as employee training mechanisms involve simulated attacks to raise awareness and educate employees about the risks and consequences of falling victim to such attacks. By conducting mock phishing exercises, organizations aim to enhance threat awareness, protect sensitive information, mitigate social engineering risks, promote incident response, and strengthen security practices. Phishing attacks are prevalent and pose significant risks to many industries, making it essential for employees to know how to defend against them. Phishing is an effective attack vector due to its exploitation of human vulnerabilities, deceptive impersonation of trusted entities, psychological manipulation, broad reach, ease of use, dynamic capabilities, adaptability, and the potential for significant financial gain. These factors make phishing attacks difficult to detect, mitigate and underscore the importance of practicing vigilance and regularly training employees to recognize and respond effectively to phishing attempts.

Through training, employees become more aware of common phishing techniques and deceptive tactics used by cybercriminals. This knowledge helps them to identify and report suspicious emails or messages, reducing the likelihood of data breaches and unauthorized access to sensitive information. By training employees to recognize phishing attempts, organizations mitigate social engineering risks. Employees learn to identify messages that use common tactics such as urgent requests, spoofed identities, and enticing offers to manipulate individuals. This knowledge helps protect employees and their organization from disclosing credentials or confidential data, or installing malware. Effective training enables employees to respond appropriately to phishing attempts, such as reporting incidents to specific IT or security teams, refraining from clicking suspicious links or opening attachments, and verifying requests sent via email using alternative channels. Training employees to recognize and respond to

phishing attempts strengthens an organization's cybersecurity defenses. It cultivates a culture of security awareness, empowers employees to protect sensitive information actively, and enhances the organization's resilience against evolving threats. Complemented by simulated phishing campaigns, regular training programs help build a knowledgeable and security-conscious workforce.

Security Awareness Training Lifecycle

Security awareness training practices typically follow a lifecycle approach consisting of several stages.

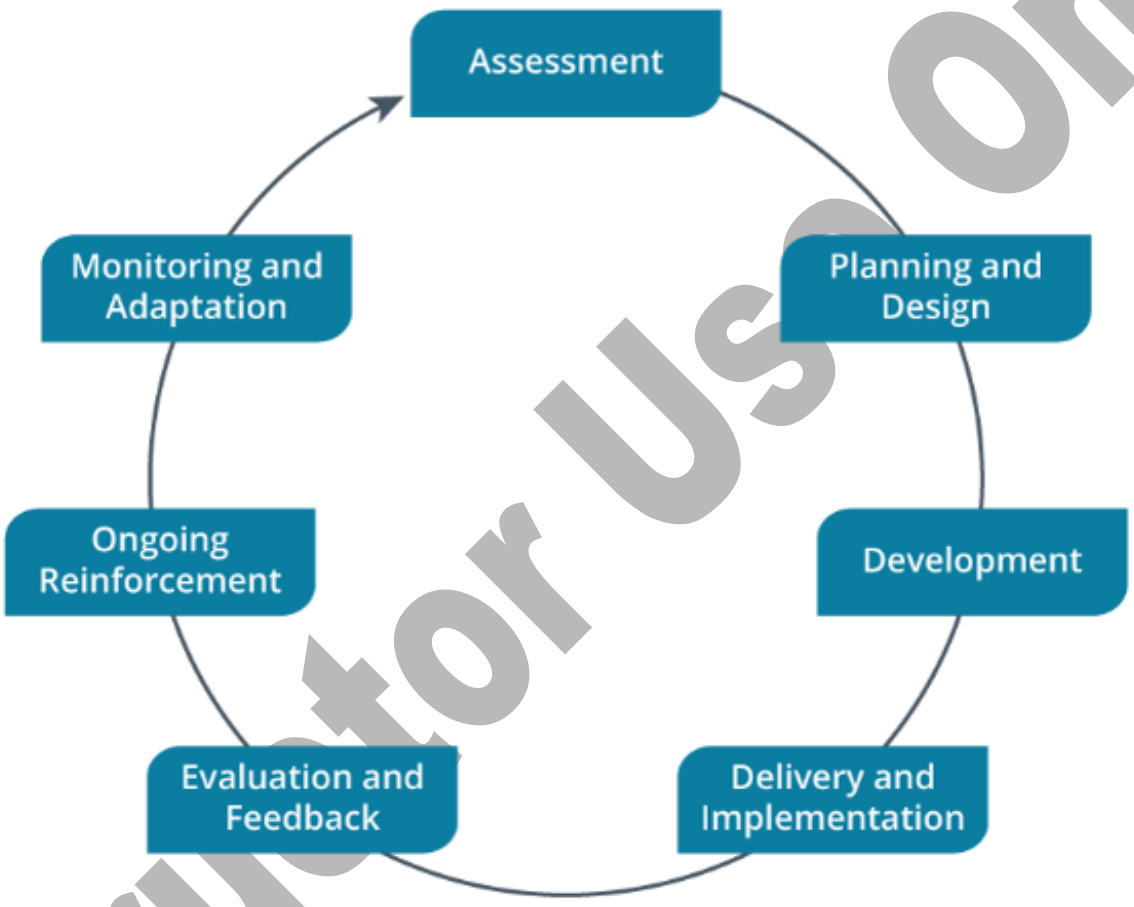

Security Awareness Training Lifecycle.

The first phase is assessing the organization's security needs and risks. Planning and designing awareness training activities follow, where a comprehensive plan is developed, including objectives, topics, and delivery methods. Once a plan is created, the development stage focuses on creating engaging and informative training materials. Training is then delivered through previously identified delivery methods, such as in-person or computer-based sessions. Evaluation and feedback activities assess the training's effectiveness and gather participant insights. Security awareness is reinforced via reoccurring training activities to ensure it remains a priority and often includes refresher training, reminders, newsletters, and awareness campaigns. Monitoring and adaptation allow organizations to continually evaluate the program's impact and make necessary adjustments based on emerging risks and changing requirements. Organizations can establish a continuous and effective security awareness training program by following this lifecycle. It helps enhance employee knowledge, steer behaviors, and cultivate a security culture within the organization. Regular assessment, evaluation, and adaptation ensure that training remains relevant and addresses evolving security threats. A well-structured security awareness training program significantly contributes to mitigating risks, protecting sensitive data, and building a resilient cybersecurity posture.

Development and Execution of Training

Successfully executing security training means effectively providing education and instruction to employees and staff that enhance their knowledge, skills, and awareness of security practices. It involves delivering training programs addressing relevant security topics, like data protection, incident response, phishing awareness, secure coding practices, physical security, and many others. Content development focuses on creating engaging and informative training materials, using clear language, and incorporating real-world examples to enhance relevance and mix interactive elements like quizzes, case studies, or simulations to encourage active participation, critical thinking, and practical application of knowledge. Facilitating dialogue, discussion, and question-and-answer sessions further enhance the learning experience. To assess the effectiveness of security awareness practices, collecting feedback, conducting assessments, and developing relevant measurements and metrics help gauge the impact of the training and identify areas for improvement. Regular reviews and updates to training materials ensure that its content remains relevant and aligned with evolving security threats. Incorporating emerging best practices and industry trends help organizations stay current and enhance their security awareness practices.

Reporting and Monitoring

One way to gauge the efficacy of security awareness training is to assess its influence initially and through ongoing evaluations.

Initial effectiveness refers to the immediate impact of security awareness training on participants. It measures the knowledge gained, awareness raised, and behavioral changes observed immediately after completing a training program. Evaluation methods can include pre-and post-training assessments, quizzes, and surveys designed to gauge participant understanding of security concepts before and after training. Measuring initial effectiveness provides insight into the immediate impact of training and how participants have absorbed the information and concepts presented.

Recurring effectiveness assesses the long-term impact and sustainability of security awareness training by examining whether participants have retained and applied the knowledge and skills gained from training in their day-to-day activities. The focus is to measure the continued behavioral changes and the level of security consciousness within the organization over an extended period.

Initial and recurring effectiveness measurements are crucial to gauge the overall impact of security awareness training. While measuring initial effectiveness shows the immediate outcomes and knowledge uptake, recurring effectiveness measurements ensure that the training has a lasting effect and leads to sustained improvements in security practices.

Assessments and Quizzes — Conducting pre- and post-training assessments and quizzes allow organizations to measure the knowledge gained by employees during training. These reports provide quantitative data regarding training effectiveness related to knowledge retention and comprehension.

Incident Reporting —Organizations can track and analyze incident reports to assess the training program's impact on incident detection and response, identifying any patterns or trends.

Phishing Simulations —Conducting simulated phishing campaigns helps organizations evaluate employees' ability to recognize and respond to phishing attempts. Reports generated from these simulations provide data on click rates, successful phish captures, and trends in susceptibility, indicating the effectiveness of the training in mitigating phishing risks.

Observations and Feedback —Managers and supervisors can provide feedback on employees' security practices and behaviors. Qualitative information such as this provides valuable insights into the practical application of training and any challenges employees face in implementing the knowledge gained.

Metrics and Performance Indicators —Tracking relevant metrics, such as the number of reported incidents, employee compliance with security policies, or changes in password hygiene, provides quantitative data on the impact of security awareness training. These metrics help measure the effectiveness of the training program over time.

Training Completion Rates —Monitoring the completion rates of security awareness training modules or sessions indicates employee engagement and adherence to training requirements. Higher completion rates suggest better participation and performance of the training content.

Anomalous Behavior

Anomalous behavior recognition refers to actions or patterns that deviate significantly from expectations. Examples include unusual network traffic, user account activity anomalies, insider threat actions, abnormal system events, and fraudulent transactions. Techniques such as network intrusion detection, user behavior analytics, system log analysis, and fraud detection are utilized to identify anomalous behavior. These techniques require monitoring and analyzing different data sources, comparing observed behavior against established baselines, and utilizing machine learning algorithms to detect deviations.

Recognizing Risky Behaviors

Risky behaviors are actions or practices that threaten data security, systems, or networks. These behaviors may involve unsafe online activities, such as clicking on suspicious links, visiting untrusted websites, or downloading unauthorized software. Risky behaviors can also include neglecting security measures, such as using weak passwords, sharing credentials, or ignoring software updates. Unexpected behaviors are actions that deviate from established security protocols or violate security policies. These behaviors can occur due to a lack of awareness, carelessness, or a failure to follow established procedures. Examples include unauthorized access to sensitive information, bypassing security controls, or disregarding physical security measures. Unintentional behaviors refer to actions without malicious intent but can still have detrimental consequences. These behaviors often stem from human error, lack of training, or lack of understanding of security best practices. Examples include accidental data breaches, mishandling of confidential information, or falling victim to social engineering attacks.

All three types of behaviors (risky, unexpected, and unintentional) can lead to security incidents, data breaches, or the compromise of sensitive information. Individuals must be aware of these behaviors, follow security guidelines, stay informed about emerging threats, and practice good cybersecurity hygiene. Organizations are responsible for training and educating employees about these behaviors to promote a security-conscious culture and minimize the impact of human-related vulnerabilities in the cybersecurity landscape.

13.2.3 Data Protection and Policies (Lesson Video)

Transcript:

In this lesson, I'm going to talk about data protection and policies. Data is a critical asset to all organizations, which must put specific policies in place for data creation, storage, and sharing.

One important data protection aspect is data retention. Data retention policies define how to organize, maintain, and archive data. They also outline how and when to dispose of information. Different data types must be legally retained for different lengths of time, depending on relevant regulations. Most data retention policies also address how to handle information involved in litigation. When you're creating a data retention policy, you need to identify the various data types your organization uses. Develop a policy that defines how long each data type should be retained as well as how they should be destroyed when the retention time expires. It's important to record this information in a policy that's clearly written. Having a written policy and ensuring that everyone in the organization follows it helps to protect you from evidence tampering accusations.
Adhering to your data retention policy protects both you and your organization. You should never allow selective or arbitrary information destruction, which could make it appear as though you're trying to hide evidence and could lead to criminal charges. Never destroy information after it's been subpoenaed or if you have reason to believe that it may be subpoenaed in the future. Evidence destruction and obstruction of justice are serious crimes that can result in imprisonment.
With this in mind, let's take a look at some examples of data retention rules. These rules aren't necessarily the ones that you'll use, but I'll go over them to give you ideas on how you can formulate your own policies.
Email inboxes can pile up quickly, so the first rule could state that all email messages are to be deleted after 90 days.
The IRS defines guidelines for maintaining tax records. To ensure compliance, you could say that your policy is to keep company tax records for seven years and employee tax records for four. After that time, they're to be shredded.

When it comes to employee HR records, a good rule of thumb is to keep them for three to four years.
You might see fit to hold onto all employee-created research and design documents to prove their ownership. So for the next rule, let's say the policy is to keep any integral research, design, or patent documents for 25 years.
Next, are vendor contracts, service-level agreements, and other types of contractual documents. A good practice is to keep these until five years after a contract has ended.
After you create your written data retention policy, consider using information classification labels to identify which retention policy rules apply to which pieces of data. This allows you to automate the data retention and destruction process. Without classification labels, it's an extremely tedious process to try to manage decades-long data records.
Having a solid data retention policy has a lot of benefits. First, it reduces the discovery-request cost in the event of legal action against your organization. Responding to discovery requests can be time consuming and costly. You significantly minimize your discovery costs if all your old, irrelevant material has already been destroyed. Second, it reduces exposure during the discovery process. By minimizing the amount of electronic material that your organization keeps on hand, you also reduce the amount of information that could potentially expose your organization to litigation. Third, it lowers the hardware and software requirements for storing old data and minimizes clutter.
That's it for this lesson. In this lesson, we talked about data protection and how to create data retention policies for your organization. We also talked about what specifically a data retention policy is and the good practices you should follow when creating one.

13.2.4 Data Protection and Policies Facts

Data protection and compliance encompass a range of practices and principles aimed at safeguarding sensitive information, ensuring privacy, and adhering to applicable laws and regulations. Data protection involves implementing measures to secure data against unauthorized access, loss, or misuse. It includes practices such as encryption, access controls, data backup, and secure storage. Compliance refers to conforming to legal, regulatory, and industry requirements relevant to data handling, privacy, security, and transparency. Organizations can safeguard individuals' privacy, ensure data security, fulfill legal requirements, and establish credibility with customers, partners, and regulatory authorities by comprehending and implementing these data protection and compliance principles. Compliance with applicable data protection laws, regulations, and standards is crucial for organizations to avoid legal liabilities, reputational damage, and financial penalties associated with noncompliance.

This lesson covers the following topics:

- Data protection
- Data loss prevention
- Data sovereignty and geographical considerations
- Data retention policies

Data Protection

Classifying data as "at rest," "in motion," and "in use" is crucial for effective data protection and security measures. By analyzing data based on its state (at rest, in motion, in use), organizations can tailor security measures and controls to address the specific risks and requirements associated with each data state. This classification helps organizations identify vulnerabilities, prioritize security investments, and ensure appropriate safeguards to protect sensitive data throughout its lifecycle. It also facilitates compliance with data protection regulations and industry best practices.

- **Data at rest** — this state means that the data is in some sort of persistent storage media. Examples of types of data that may be at rest include financial information stored in databases, archived audiovisual media, operational policies and other management documents, system configuration data, and more. In this state, it is usually possible to encrypt the data using techniques such as whole disk encryption, database encryption , and file- or folder-level encryption. It is also possible to apply permissions—access control lists (ACLs)—to ensure only authorized users can read or modify the data. ACLs can be applied only if access to the data is fully mediated through a trusted OS.

Data in transit (or data in motion) — this is the state when data is transmitted over a network. Examples of types of data that may be in transit include website traffic, remote access traffic, data being synchronized between cloud repositories, and more. In this state, data can be protected by a transport encryption protocol, such as TLS or IPSec.

With data at rest, there is a greater encryption challenge than with data in transit, as the encryption keys must be kept secure for longer. Transport encryption can use ephemeral (session) keys.

Data in use (or data in processing) — this is the state when data is present in volatile memory, such as system RAM or CPU registers and cache. Examples of types of data that may be in use include documents open in a word processing application, database data that is currently being modified, event logs being generated while an operating system is running, and more. When a user works with data, it usually needs to be decrypted as it goes from in rest to in use. The data may stay decrypted for an entire work session, which puts it at risk. However, trusted execution environment (TEE) mechanisms, such as Intel Software Guard Extensions (software.intel.com/content/www/us/en/develop/topics/software-guard-extensions/details.html), can encrypt data as it exists in memory so that an untrusted process cannot decode the information.

Due diligence in the context of data protection describes the comprehensive assessment and evaluation of an organization's data protection practices and measures. It involves examining and verifying the adequacy of data security controls, privacy policies, data handling procedures, and compliance with applicable laws and regulations.

Data Loss Prevention

To apply data guardianship policies and procedures, smaller organizations might classify and type data manually. However, an organization that creates and collects large amounts of personal data will usually need to use automated tools to assist with this task. There may also be a requirement to protect valuable intellectual property (IP) data. Data loss prevention (DLP) products automate the discovery and classification of data types and enforce rules so that data is not viewed or transferred without proper authorization. Such solutions will usually consist of the following components:

Policy server — to configure classification, confidentiality, and privacy rules and policies, log incidents, and compile reports.

Endpoint agents — to enforce policy on client computers, even when they are not connected to the network.

Network agents — to scan communications at network borders and interface with web and messaging servers to enforce policy.

DLP agents scan content in structured formats, such as a database with a formal access control model, or unstructured formats, such as email or word processing documents. A file cracking process is applied to unstructured data to render it in a consistent scannable format. The transfer of content to removable media, such as USB devices, by email, instant messaging, or even social media, can be blocked if it does not conform to a predefined policy. Most DLP solutions can extend the protection mechanisms to cloud storage services, using either a proxy to mediate access or the cloud service provider's API to perform scanning and policy enforcement.

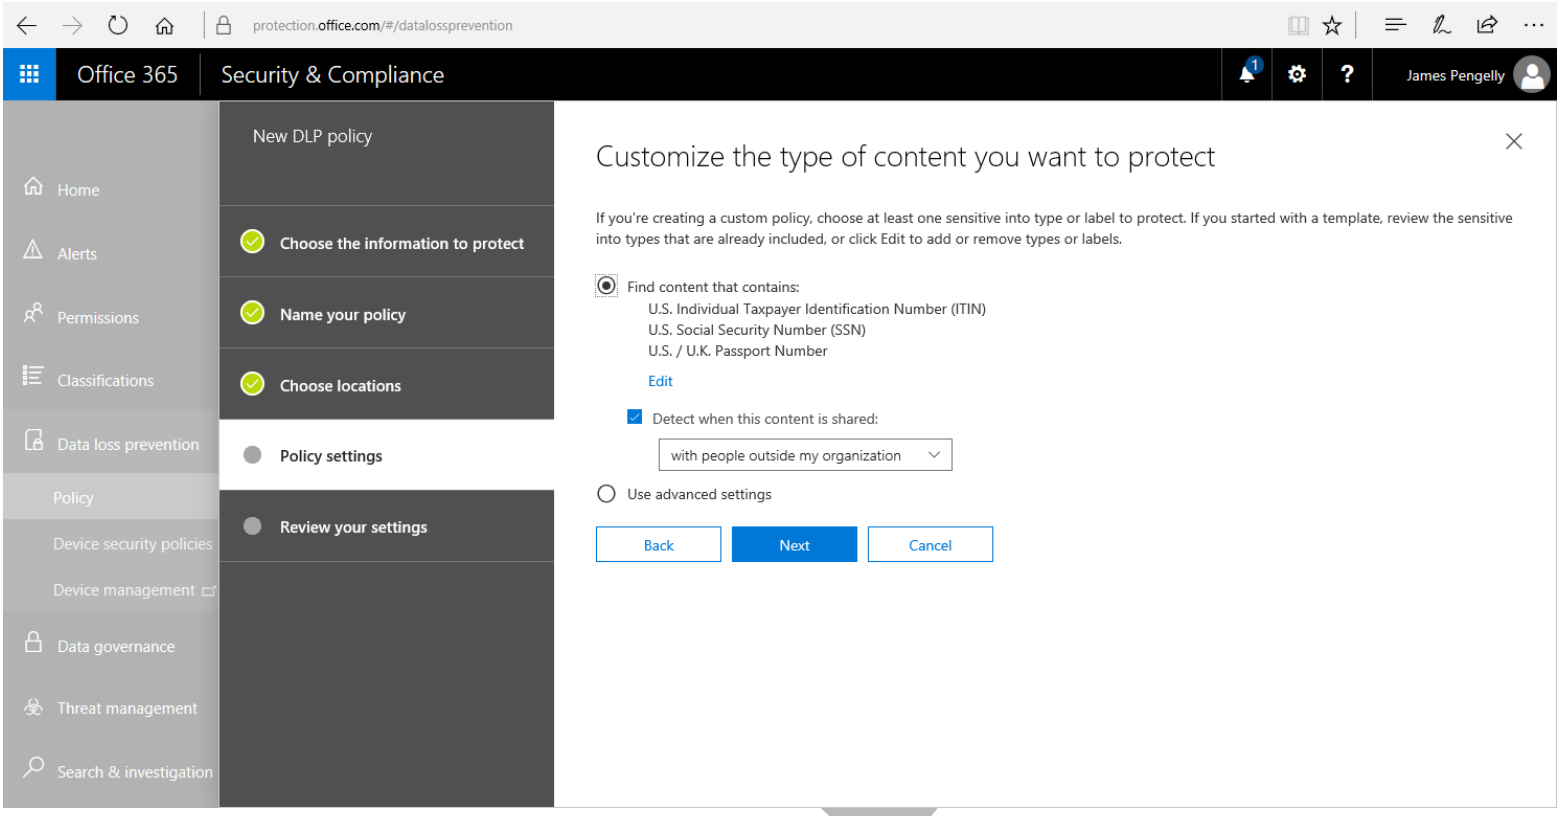

Creating a DLP policy in Office 365. (Screenshot used with permission from Microsoft.)

Remediation is the action the DLP software takes when it detects a policy violation. The following remediation mechanisms are typical:

Alert only — copying is allowed, but the management system records an incident and may alert an administrator.

Block — the user is prevented from copying the original file but retains access. The user may or may not be alerted to the policy violation, but it will be logged as an incident by the management engine.

Quarantine — access to the original file is denied to the user (or possibly any user). This might be accomplished by encrypting the file in place or by moving it to a quarantine area in the file system.

Tombstone — the original file is quarantined and replaced with one describing the policy violation and how the user can release it again.

When configured to protect a communications channel such as email, DLP remediation might take place using client-side or server-side mechanisms. For example, some DLP solutions prevent the actual attaching of files to the email before it is sent. Others might scan the email attachments and message contents and then strip out certain data or stop the email from reaching its destination.

Data Sovereignty and Geographical Considerations

Some states and nations may respect data privacy more or less than others, and likewise, some nations may disapprove of the nature and content of certain data. They may even be suspicious of security measures such as encryption. When your data is stored or transmitted in other jurisdictions or when you collect data from citizens in other states or other countries, you may not "own" the data in the same way as you'd expect or like to.

Data Sovereignty

Data sovereignty refers to a jurisdiction preventing or restricting processing and storage from taking place on systems that do not physically reside within that jurisdiction. Data sovereignty may demand certain concessions on your part, such as using location-specific storage facilities in a cloud service.

For example, GDPR protections are extended to any EU citizen while they are within EU or EEA (European Economic Area) borders. Data subjects can consent to allow a transfer, but there must be a meaningful option for them to refuse consent. If the transfer destination jurisdiction does not provide adequate privacy regulations (to a level comparable to GDPR), then contractual safeguards must be given to extend GDPR rights to the data subject. In the United States, companies can self-certify that the protections they offer are adequate under the Privacy Shield scheme (privacyshield.gov/US-Businesses).

Maintaining compliance with data sovereignty requirements requires several approaches. Organizations ensure data localization by storing and processing data using datacenters or cloud providers within defined legal or geographic boundaries. Additionally, contractual agreements with vendors and service providers ensure data remains within approved boundaries by outlining responsibilities, restrictions, and mandatory safeguards.

Geographical Considerations

Geographic access requirements fall into two different scenarios:

- Storage locations might have to be carefully selected to mitigate data sovereignty issues. Most cloud providers allow a choice of datacenters for processing and storage, ensuring that information is not illegally transferred from a particular privacy jurisdiction without consent.

- Employees needing access from multiple geographic locations. Cloud-based file and database services can apply constraint-based access controls to validate the user's geographic location before authorizing access.

Geographic restrictions impact other business functions:

- Geolocation requirements impact data protection practices by requiring organizations to ensure data remains within a designated boundary, such as utilizing local datacenters or cloud providers. Geolocation restrictions affect data protection practices such as data replication and data dispersion.

- Geolocation requirements impact incident investigation and forensics activities because they often include jurisdiction-specific data access and sharing restrictions and other legal requirements.

Data Retention Policies

Data retention policies define how information in your possession is maintained and for how long. The key point to remember is that different types of data must be retained for different lengths of time based on legal and business requirements.

Data retention policies also typically describe procedures for:

- Archiving information

- Destroying information when the retention limit is reached

- Handling information involved in litigation

Review all the different types of information used in your organization and develop a policy that defines how long different types of data are retained and destroyed when the retention period is passed. Record this information in a clearly written policy, and never

allow selective or arbitrary information destruction. It might make it appear that you are trying to hide evidence and could expose you to potential criminal charges.

Benefits of Data Retention Policies

Having a written policy and ensuring everyone in the organization follows it protects you from accusations of destroying evidence. Adhering to your data retention and destruction policy protects you and your organization.

Other benefits of implementing a data retention policy include:

- Reduced cost of discovery requests in the event of legal action. Responding to discovery requests can be time-consuming and costly. If old material has been destroyed, discovery costs are minimized.
- Reduced exposure during discovery. Minimizing the amount of electronic material an organization keeps reduces the amount of information that could expose an organization to potential litigation.
- Reduced hardware and software requirements for storing old data.

Data Retention Best Practices

Some best practices include:

- Delete email messages after 90 days.
- Keep tax-related information for seven years. This timeframe should be defined by the applicable taxation authority. For example, the United States Internal Revenue Service requires tax information to be retained for seven years.
- Keep employee records for four years after an employee leaves the organization.
- Keep integral research, design, or patent documents for 25 years.
- Keep contracts with vendors and partners for five years after a contract has ended.
- Delete employee files after one year.
- All information should be destroyed before being disposed of. Simply deleting files can leave sensitive information behind.
- After creating your written data retention policy, use information classification labels to identify which retention policy rule is to be applied to specific data. Using classification labels allows you to use software tools to automate the data retention and destruction process.

13.2.5 Practice Questions (Section Quiz)

q_pers_pol_aup_01_secp8

A medium-sized organization is undergoing an audit for its information security practices. As a security analyst, the auditor seeks to assess the organization's use of an acceptable use policy (AUP).

What crucial aspect of the AUP should the auditor focus on to ensure the organization meets the standards set for information security?

Answers:

***The AUP includes clear consequences for noncompliance.**

The AUP includes guidance for personal use of organizational resources.

The AUP includes the number of allowed password attempts before locking an account.

The AUP includes a list of approved software for each department.

Explanation:

Any acceptable use policy crucially outlines clear consequences for noncompliance, which serves to deter potential violators and provides a clear path for action when violations occur.

While an AUP might offer guidance for personal use of organizational resources, its primary focus remains on regulating the use of these resources to maintain security.

A password policy or an account lockout policy, not an AUP, typically includes the provision for the number of allowed password attempts before locking an account.

The policy known as the approved software list, not the AUP, includes the list of approved software for each department.

q_pers_pol_aup_02_secp8

An organization recently experienced a security breach due to the actions of an employee who engaged in an activity that posed a risk to the company's information systems. The employee downloaded unverified software onto the company device, resulting in a malware infection.

Following this incident, the company plans to implement a policy to prevent similar occurrences in the future.

Which of the following policies is MOST suitable for addressing this specific issue?

Answers:

Business continuity and continuity of operations plans (COOP)

***Acceptable use policy (AUP)**

Disaster recovery policy

Change management policy

Explanation:

AUP sets the standard for acceptable behavior by users on network and computer systems. The AUP typically comprises rules about software downloads, aiming to deter users from participating in activities that could damage the organization or its resources.

COOP policies strive to keep critical processes running during and following significant disruption. However, these policies do not directly tackle the problem of an employee downloading risky software.

A disaster recovery policy concentrates on the resumption of operations after a catastrophic event but does not incorporate preventive measures against risky behaviors by users, such as unauthorized software downloads.

A change management policy describes the process for requesting, reviewing, approving, and implementing changes to IT systems and software.

q_pers_pol_aup_03_secp8

A new employee reviews key guidance given by the department manager. The manager reminds the new employee to pay close attention to the directive that concerns all electronic usage.

What documents should the employee review? (Select two.)

Answers:

***AUP**

***Information security policy**

Change management

SDLC

Clean desk policy

Explanation:

An acceptable use policy (AUP) outlines how individuals may use network and computer systems by defining what constitutes acceptable behavior by users. AUPs typically address browsing behavior, appropriate content, software downloads, and handling sensitive information.

Information security policies are policies an organization creates to ensure that all information technology users comply with rules and guidelines related to the security of the information stored within the environment or the organization's sphere of authority.

Change management policies govern the requests, reviews, approvals, and implementation of IT systems and software changes.

Software Development Lifecycle (SDLC) policies govern software development within an organization. They do not govern what employees can do with electronics.

A clean desk policy means that each employee's work area should be free from any documents left there. The aim of the policy is to prevent sensitive information from being obtained by unauthorized staff or guests at the workplace. It does not focus on electronic usage.

q_pers_pol_clean_desk_01_secp8

An organization observes employees leaving sensitive documents on their desks, thereby exposing sensitive data in the work area. To stop unauthorized staff or guests from accessing this information, the organization decides to introduce a new policy.

Which policy would resolve this issue?

Answers:

Acceptable use policy

***Clean desk policy**

Code of conduct and social media analysis

Use of personally owned devices in the workplace

Explanation:

A clean desk policy ensures that each employee's work area is free from any documents left there, preventing unauthorized staff or guests from obtaining sensitive information.

Acceptable use policy mainly concerns using the organization's equipment and technology, forbidding actions such as defrauding, defaming, obtaining illegal material, and unauthorized access to confidential data.

Code of conduct and social media analysis do not directly address the issue of sensitive documents left unattended on desks.

This policy concerns the threats portable devices pose to data security, including file copying and recording functions and the unauthorized use of personal software. It does not address the issue of sensitive documents left out in the open.

q_pers_pol_clean_desk_02_secp8

A company had data for an upcoming project stolen and leaked online. The investigation implies that social engineering is the cause.

Which policy can prevent such an incident from occurring?

Answers:

***Clean desk**

Fair use

Standard operating procedure (SOP)

Non-disclosure agreement (NDA)

Explanation:

A clean desk policy means an employee's work area should be free from documents or information. The policy aims to prevent sensitive information from prying eyes.

A fair use (or acceptable use) policy defines what someone can do while using a particular service or resource (such as using systems at a business).

A standard operating procedure (SOP) is a documented list of steps or actions to perform a task to a specified and expected standard.

A non-disclosure agreement (NDA) is a document between parties (such as an employer and employee) for a legal basis in protecting information assets. In this situation the information was stolen, not disclosed by an employee, and social engineering was involved.

q_pers_pol_clean_desk_03_secp8

A user invited a friend to the office for a tour. There were several moments when no one was around the visitor. During these times, the friend might have viewed confidential information.

What type of policy can prevent such an activity from occurring?

Answers:

***Clean desk**

Fair use

Standard operating procedure (SOP)

Non-disclosure agreement (NDA)

Explanation:

A clean desk policy means an employee's work area should be free from any documents or information. The policy aims to prevent sensitive information from prying eyes.

A fair use (or acceptable use) policy defines what someone can do while using a particular service or resource (such as using systems at a business).

A standard operating procedure (SOP) is a documented list of steps or actions to perform a task to a specified and expected standard.

A non-disclosure agreement (NDA) is put into place between parties (such as an employer and employee) as a legal basis for protecting information assets.

q_pers_pol_incident_response_secp8

Which policy outlines the processes to follow after a security breach or cyberattack occurs and includes procedures for identifying, investigating, controlling, and mitigating the impact of incidents?

Answers:

***Incident response policy**

Acceptable use policy (AUP)

Disaster recovery policy

Change management policy

Explanation:

The incident response policy is vital in managing cybersecurity incidents effectively. It outlines the steps to follow after a security breach or cyberattack, providing guidelines for identifying, investigating, controlling, and mitigating the impact of such incidents.

The acceptable use policy (AUP) outlines acceptable uses for network and computer systems, defining appropriate user behavior and addressing consequences for noncompliance.

The disaster recovery policy details steps required to recover from catastrophic events like natural disasters or significant security breaches, aiming to restore operations quickly and efficiently.

A change management policy governs the request, review, approval, and implementation of IT systems and software changes, including documentation requirements.

q_pers_pol_offboarding_secp8

An organization has hired an HR director to improve the performance of the HR division. The director first noted the lack of an offboarding process for employees or contractors.

What are some IT security areas an offboarding process should focus on? (Select three.)

Answers:

***Account management**

***Physical security**

Personal assets

***Company assets**

Clean desk

Acceptable use

Asset allocation

Explanation:

The following are focus areas during the offboarding process:

When a user leaves, the HR division immediately deactivates the account, marking it for auto-deletion later.

While not always related to IT security, physical security and the items that support physical security are important to keep track of.

An organization must account for any electronic devices employees use during employment, returning them before their departure. An inventory sheet allows for the accounting of all of the property previously in their control.

The following are not the direct focus of an offboarding process:

Any personal devices (assets) used for company activities must have the proprietary information removed, including the removal of any software purchased by the company.

A clean desk means that each employee's work area should be free from any documents left there, which is not a primary focus of offboarding.

An acceptable use policy (AUP) is important to protect the organization from the security and legal implications of employees misusing its equipment. However, it is not a primary focus of offboarding.

Asset allocation provisions computers or mobile devices for the user, including the use of bring-your-own-device handsets. This is an onboarding issue, not an offboarding issue.

q_pers_pol_phishing_sims_secp8

As the head of the IT department in your organization, you have noticed an increase in the number of phishing attempts targeting your employees. You have implemented several measures to combat this, but the attempts continue to rise.

Which of the following strategies would be the most effective in addressing this issue?

Answers:

Implement a stricter internet usage policy and block access to all non-work-related websites.

***Conduct regular phishing simulations to test employee awareness and response.**

Increase the frequency of system-wide password changes.

Install more advanced antivirus software on all company computers.

Explanation:

Conducting regular phishing simulations is the most effective strategy. This approach not only tests the current level of employee awareness but also provides a practical, hands-on learning experience for employees. It allows you to identify who is most at risk and tailor your training accordingly. This method directly addresses the issue and helps to improve the overall security posture of the organization.

Implementing a stricter internet usage policy and blocking access to all non-work-related websites may reduce the risk of phishing attempts, but it could also hinder employee productivity and morale. It does not directly address the issue of phishing and does not provide employees with the knowledge or skills to identify and avoid phishing attempts.

Increasing the frequency of system-wide password changes can improve security, but it does not directly address phishing attempts. Furthermore, it could lead to employee frustration and potentially weaker passwords if employees struggle to remember their constantly changing login credentials.

Installing more advanced antivirus software can help to detect and block phishing attempts, but it is not a foolproof solution. Some phishing attempts may still get through, and without proper awareness and training, employees may still fall for these attempts.

q_pers_pol_risky_behaviors_secp8

As a cybersecurity analyst, you are tasked with identifying risky behaviors among employees that could potentially lead to security breaches.

Which of the following behaviors would you consider to be the MOST risky and why?

Answers:

An employee frequently works late and accesses the company network after hours.

An employee uses their personal smartphone to check work emails.

***An employee regularly uses the same password for multiple systems and applications.**

An employee occasionally forgets to log out of their workstation when they leave for lunch.

Explanation:

Using the same password for multiple systems and applications is the most risky behavior. If one system or application is compromised, all others that use the same password are also at risk. This behavior makes it easy for attackers to gain access to multiple areas of the network, potentially leading to a significant security breach.

While an employee accessing the company network after hours could potentially pose a risk, it is not inherently risky behavior. Many employees work late or outside of normal business hours. The key is to ensure that the network is secure at all times, regardless of when it is being accessed.

Using a personal smartphone to check work emails can pose a risk if the phone is lost or stolen, or if it is infected with malware. However, with proper security measures in place, such as encryption and remote wipe capabilities, this risk can be mitigated.

Forgetting to log out of a workstation can pose a risk, particularly if the workstation is left unattended in a public or shared space. However, this risk can be mitigated with automatic log-out settings and screen locks. While it is a risky behavior, it is not as risky as using the same password for multiple systems and applications.

q_pers_pol_training_01_secp8

A software development company recognizes that some of its employees are vulnerable to phishing attacks. To address this, the company plans to set up a training program.

What factors should the company primarily consider while defining such training programs?

Answers:

The job titles of the employees

The personal interests of the employees

***The roles performed by the employees**

The length of service of the employees

Explanation:

Employees may perform different roles and have different security training, education, or awareness requirements in each role. Therefore, focusing on job roles rather than job titles is essential in setting up effective security training programs.

Job titles might not accurately reflect an employee's roles, responsibilities, and security training needs.

While it may be useful to consider personal interests to enhance engagement, this should not be the primary factor in defining security training programs.

While the length of service might influence an employee's experience or familiarity with security issues, it is not a primary factor in determining the training needs for specific roles.

q_pers_pol_training_02_secp8

A software application contains sensitive transmittal information, and an end user takes it out on a laptop in the field. The end user must understand how to protect and dispose of the data.

Which one of the following should help the end user prepare for this?

Answers:

***User training**

Vendor-specific guide

General purpose guide

Change management

Explanation:

User training teaches users new functionality and proper policies and procedures for both the company and the software. Users should complete training before using the system to prevent incidents and to understand what to do in the event of one.

Vendor-specific guides provide instructions on installing and securely configuring hardware and software for a specific vendor.

General purpose guides help increase security in hardware and software by providing instructions for configuring a system based on roles and appliances.

Change management is a process that involves the prevention of unauthorized changes to a system. This process protects from unwanted outages.

q_pers_pol_training_lifecycle_secp8

As the security awareness training manager at your company, you have recently completed a round of training sessions on the company's security policies.

However, a few weeks later, you notice that some employees are not adhering to the policies, particularly those related to secure email practices.

What should you do?

Answers:

Ignore the issue, assuming that employees will eventually learn from their mistakes.

Immediately fire the employees who are not following the secure email practices.

***Implement ongoing reinforcement strategies, such as regular reminders, mini-training sessions, and visual aids around the office, to keep the secure email practices fresh in employees' minds.**

Discipline the employees by restricting their access to certain company resources.

Explanation:

Ongoing reinforcement is a key part of the security awareness training lifecycle and the correct answer. Regular reminders, mini-training sessions, and visual aids can help keep the secure email practices fresh in employees' minds, leading to better adherence to the policies.

Ignoring the issue does not solve the problem and could potentially lead to more severe security breaches. As the security awareness training manager, it is your responsibility to address any security issues that arise.

Firing employees is not a constructive solution to the problem. It does not address the root cause of the issue, which is the lack of understanding or adherence to password management guidelines.

Disciplining employees by restricting their access to company resources does not solve the problem. It could also lead to a decrease in productivity and employee morale. The goal of the Security Awareness Training program is to educate employees, not to punish them.

q_pers_pol_update_training_secp8

A company has recently experienced a significant increase in phishing attacks. The IT department has implemented technical controls to block phishing emails, but some are still getting through.

The company has a security awareness training program, but it has not been updated in several years.

What should be the company's immediate action to address this issue?

Answers:

***Update the security awareness training program to include information about the latest phishing techniques and conduct a company-wide training session.**

Invest in more advanced technical controls to block phishing emails.

Conduct a company-wide training session using the existing security awareness training program.

Discipline employees who fall for phishing attacks to deter others from making the same mistake.

Explanation:

Updating the security awareness training program to include information about the latest phishing techniques and conducting a company-wide training session is the correct answer and would help employees recognize and avoid phishing emails. This is a proactive approach that addresses the root cause of the problem.

While investing in more advanced technical controls might help block more phishing emails, it doesn't address the issue of employees not being able to recognize and avoid phishing emails that do get through.

Conducting a company-wide training session using the existing security awareness training program would not be as effective because the program has not been updated in several years and likely does not include information about the latest phishing techniques.

Disciplining employees who fall for phishing attacks might create a culture of fear and discourage employees from reporting when they do fall for phishing attacks. This could potentially make the problem worse. It's more effective to educate employees so they can recognize and avoid phishing attacks.

q_dataprot_pol_adhere_secp8

As the data protection officer at a large corporation, you have just completed the development of a comprehensive data retention policy. The policy includes specific retention periods for different types of data, procedures for archiving and destroying information, and guidelines for handling information involved in litigation.

A few weeks after the policy has been implemented, you discover that a department within the corporation has been selectively deleting certain types of data before the retention limit is reached. The department head argues that the data is not relevant to their operations and is taking up unnecessary storage space.

How should you respond?

Answers:

Allow the department to continue deleting the data, but ask them to document their actions for transparency.

***Insist that the department adhere to the data retention policy, explaining the potential legal and business implications of not doing so.**

Modify the data retention policy to allow for selective deletion of data if it is deemed irrelevant to a department's operations.

Report the department's actions to senior management and recommend disciplinary action.

Explanation:

Insisting that the department adhere to the data retention policy is the correct answer. As the data protection officer, it is your responsibility to ensure that the data retention policy is adhered to across the corporation. The policy was put in place to protect the corporation and its data, and any deviations from it could have serious implications.

Allowing the department to continue deleting the data is not the best because it undermines the integrity of the data retention policy. While transparency is important, allowing selective deletion of data could expose the corporation to potential legal risks and accusations of evidence destruction.

Modifying the data retention policy to allow for selective deletion of data is not advisable because it could lead to inconsistencies in how data is handled across the corporation. It could also potentially expose the corporation to legal risks if data is deleted prematurely.

While it's important to report non-compliance, recommending disciplinary action immediately may not be the best first step. It's better to first ensure the department understands the importance of the data retention policy and the potential implications of not adhering to it.

q_dataprot_pol_data_at_rest_secp8

A system administrator implemented encryption across the organization's IT infrastructure. The infrastructure includes various types of data storage methods.

Which of the following data storage methods can the system administrator encrypt to increase the security of data at rest? (Select three.)

Answers:

***Volume**

***File**

***Partition**

User interface

Website traffic

Remote access

Cloud repository synchronization

Explanation:

A volume is a single accessible storage area spanning a single partition on a hard drive or encompassing an entire storage device. When encrypted, an attacker cannot decipher the information without the encryption key.

File-level encryption allows for the encryption of individual files as opposed to an entire disk or volume.

Similar to volume encryption, partition encryption converts the entirety of the data within the specified partition into an unreadable format without the correct encryption key.

The following are data storage methods that are not designed to increase the security of data at rest:

The user interface (UI) is how users interact with a system or software. Since the UI is not a data storage method, it does not contain data at rest and, therefore, does not require encryption.

Website traffic, remote access, and cloud repository synchronization are all associated with data in transit.

q_dataprot_pol_data_in_use_secp8

A recent security flaw allowed a malicious actor to access sensitive data even though the data never left the server and there is full drive encryption.

Which data state did the malicious actor MOST likely compromise?

Answers:

In transit

At rest

***In use**

Through bluetooth

Explanation:

Data in use (or data in processing) refers to the state in which data is present in volatile memory, such as system Random Access Memory (RAM) or Central Processing Unit (CPU) registers and cache. The security flaw allows for data exploitation while in use.

Data in transit (or data in motion) is the state in which data transmits over a network. In the scenario, the data never left the server.

Data at rest refers to the state when the data is in some persistent storage media. In the scenario, there is hard drive encryption.

Bluetooth would involve data in transit and not relate to the exposure of data in use.

q_dataprot_pol_dlp_secp8

After a recent breach, an organization mandates increased monitoring of corporate email accounts.

What can the organization use that mediates the copying of tagged data to restrict it to authorized media and services and monitors statistics for policy violations?

Answers:

***Data loss prevention**

Antivirus (A-V)

Security content automation protocol

Simple Network Management Protocol (SNMP) trap

Explanation:

Data loss prevention (DLP) mediates the copying of tagged data to restrict it to authorized media and services and monitors statistics for policy violations.

Antivirus (A-V) software detects and removes infections and (in most cases) other types of malware, such as worms, Trojans, rootkits, adware, spyware, password crackers, network mappers, and denial-of-service (DoS) tools.

Security content automation protocol (SCAP) allows compatible scanners to determine whether a computer meets a configuration baseline.

A Simple Network Management Protocol (SNMP) trap informs the management system of a notable event such as port failure, chassis overheating, power failure, or excessive Central Processing Unit (CPU) utilization.

q_dataprot_pol_due_diligence_01_secp8

A healthcare provider is preparing for an upcoming audit of its patient data management system. The chief compliance officer focuses on ensuring that the organization has taken the necessary steps to identify and minimize risks related to the handling of patient data.

What is the chief compliance officer primarily concentrating on?

Answers:

***Due diligence and care**

Attestation and acknowledgment

Data encryption strategy

Employee training program

Explanation:

Due diligence and care refer to the organization's responsibility to identify and minimize risks. In an upcoming audit for patient data management, the chief compliance officer focuses on taking necessary steps to identify and minimize risks.

Attestation and acknowledgment involve formally asserting that the organization has met certain conditions or requirements, typically through signed documents.

A data encryption strategy is a critical part of data protection. However, it is not the focus of identifying and minimizing risks, which is the primary task of the chief compliance officer.

Training is essential but is not the focal point of the chief compliance officer's current task of identifying and minimizing risks through careful examination and action.

q_dataprot_pol_due_diligence_02_secp8

A new tech firm creates measures to ensure it adheres to all compliance and data privacy aspects.

What describes the comprehensive assessment and evaluation of an organization's data protection practices and measures?

Answers:

***Due diligence**

Attestation and acknowledgment

Data inventories

Data retention

Explanation:

Due diligence in data protection describes the comprehensive assessment and evaluation of an organization's data protection practices and measures.

Attestation and acknowledgment require individuals or entities to formally acknowledge their understanding of compliance obligations and commitment to adhere to them through signed agreements, policy acknowledgments, and training activities.

Data inventories provide a comprehensive overview of the types of handled data, the purposes for processing, the legal basis, and the recipients of the data to ensure transparency and accountability.

Data retention is another area impacted by privacy laws. Organizations must retain personal data only for as long as necessary to fulfill the intended purpose or as required by law.

q_dataprot_pol_due_diligence_03_secp8

A software technician develops a new procedure to safeguard privacy data and ensure all groups adhere to compliance mandates.

What BEST describes due diligence?

Answers:

***It is the comprehensive assessment and evaluation of an organization's data protection practices.**

It requires individuals or entities to formally announce their understanding of compliance obligations.

It provides a comprehensive overview of the types of handled data.

It is an established timeline that requires organizations to keep documentation.

Explanation:

Due diligence in data protection describes the comprehensive assessment and evaluation of an organization's data protection practices and measures.

Attestation and acknowledgment require individuals or entities to formally acknowledge their understanding of compliance obligations and commitment to adhere to them through signed agreements, policy acknowledgments, and training activities.

Data inventories provide a comprehensive overview of the types of handled data, the purposes for processing, the legal basis, and the recipients of the data to ensure transparency and accountability.

Data retention is another area impacted by privacy laws. Organizations must retain personal data only for as long as necessary to fulfill the intended purpose or as required by law.

q_dataprot_pol_encryption_01_secp8

A financial organization has hired a cybersecurity expert to strengthen the security of its system.

The expert recommends implementing a specific technique into unreadable ciphertext by converting plaintext credit card information, regardless if it is active, in transit, or at rest.

What technique should the cybersecurity expert implement?

Answers:

***Encryption**

Tokenization

Obfuscation

Hashing

Explanation:

Encryption is a process that transforms plaintext information into an unreadable format using an encryption algorithm and a key. This process is suitable for protecting sensitive data in all states: at rest, in transit, or in use.

Tokenization provides a level of data protection, but it does not transform the original data into an unreadable format.

Obfuscation is a method of making something difficult to understand or find, such as a message or data, but it does not necessarily convert data into an unreadable format.

Hashing converts data into a fixed-size string of characters, representing the data's digital fingerprint. However, hashing is a one-way function. Upon hashing data, it is irreversible, and the original data is irretrievable.

q_dataprot_pol_encryption_02_secp8

A sole proprietorship construction company contacted an information technology (IT) consultant for technical support for a computer issue. After resolving that issue, the consultant suggested the construction company enable computer encryption.

Why might the company want to enable encryption on its computers' hard drives?

Answers:

To slow down data removal from a stolen device.

To prevent phishing

***To prevent data removal from a stolen device**

To prevent theft

Explanation:

Enabling hard drive encryption is a basic step to prevent data loss in the event of a stolen device. Without it, anyone can easily access the stolen device, regardless of needing a password.

The purpose of hard drive encryption is to prevent data loss in the event of a stolen device. However, as technology improves, malicious actors can exploit weaknesses in encryption.

Hard drive encryption does not prevent phishing. However, malicious actors could use information from an unencrypted hard drive to phish other people.

Hard drive encryption may dissuade theft, but it will not prevent theft. A malicious actor can still wipe a hard drive to reuse it.

q_dataprot_pol_encryption_03_secp8

A chief executive officer pushed back against the information technology department's proposal to set up disk encryption on all devices.

What BEST describes why the CEO should approve the proposal instead of pushing back against it?

Answers:

***Disk encryption protects the company from data theft in case of stolen devices.**

Disk encryption slows down a computer's performance.

The costs are not worth incurring.

The company does not have enough sensitive data.

Explanation:

Disk encryption protects against data theft when a malicious actor steals a device. The data remains safe as long as the malicious actor does not have the keys.

Disk encryption will slow performance and can be a negative impact. However, most modern disk encryption suites only lock down access to the data when the device is off.

Disk encryption can save a company significant costs associated with loss. The cost of setting up disk encryption is significantly less than the financial impacts on a company.

Even a limited amount of sensitive data is worth protecting with encryption to save financial and legal costs associated with loss or theft of the data.

q_dataprot_pol_geolocation_secp8

A cloud storage company wants to ensure that it stores user data in the same country as the user's physical location.

Which factor is the cloud storage company primarily considering here?

Answers:

***Geolocation**

Data sovereignty

Data classification

Data integrity

Explanation:

Geolocation refers to the identification of the geographic location of an object, such as a computing device or server.

Data sovereignty is relevant in this context. However, in this scenario, the primary factor the cloud storage company is considering is where it should store user data based on the user's location, which is a task for geolocation.

Data classification involves categorizing data based on its sensitivity level (i.e., public, confidential). This process can help organizations determine appropriate security measures, but it is not the primary factor the cloud storage company is considering.

Data integrity is maintaining and assuring the accuracy and consistency of data over its lifecycle. While a critical aspect of data management, it is not the primary consideration.

q_dataprot_pol_keep_emails_secp8

You are the IT manager at a mid-sized company. The company has a data retention policy in place that requires all emails to be deleted after 90 days.

However, the sales Ddpartment has been keeping all their emails for potential customer follow-ups and references. They argue that these emails are crucial for their operations and customer relations.

How should you handle this situation?

Answers:

Allow the sales department to keep their emails indefinitely as they are crucial for their operations.

Insist that the sales department adheres to the company's data retention policy and deletes all emails after 90 days.

***Work with the sales department to identify a solution that allows them to keep important customer-related emails while still adhering to the data retention policy.**

Report the sales department's non-compliance to the company's senior management and let them handle the situation.

Explanation:

Working with the sales department to identify a solution is the correct answer. It's important to maintain the integrity of the data retention policy while also considering the unique needs of different departments. Working with the sales department to identify a solution that allows them to keep important customer-related emails could be a win-win situation.

Allowing the sales department to keep their emails indefinitely is not advisable because it undermines the company's data retention policy. Allowing one department to keep their emails indefinitely could lead to inconsistencies and potential legal risks.

Insisting that the sales department adheres to the company's data retention policy, while adhering strictly to the policy, does not take into consideration the unique needs of the sales department. It could potentially harm the company's customer relations and sales operations.

While it's important to report non-compliance, it's also crucial to work towards a solution first. Reporting the situation to senior management without trying to resolve it first could lead to unnecessary conflicts and misunderstandings.

q_dataprot_pol_permission_01_secp8

A systems administrator receives an alert for potential unauthorized access to sensitive data while in active memory on a server within the organization. The organization has tasked the systems administrator with enforcing stricter controls to prevent such breaches.

What would be the MOST appropriate measure to implement?

Answers:

***Permission restrictions**

Data encryption

Data masking

Data obfuscation

Explanation:

Permission restrictions are an ideal option to secure data in use. Implementing permission restrictions would limit access to only authorized users, reducing the likelihood of unauthorized access.

Data encryption is an effective way to protect data at rest or in transit, but it does not typically secure data while the system processes it or when it is in active memory.

Data masking creates a structurally similar but inauthentic version of an organization's data and is not the most appropriate method for securing data in use.

Data obfuscation makes data harder to understand, making it difficult for an unauthorized user to make sense of the data. However, it is less appropriate for data in active memory compared to permission restrictions.

q_dataprot_pol_permission_02_secp8

A network engineer in an organization receives a notification regarding a potential data breach involving confidential information currently in active memory on one of their servers. The organization has decided to enforce stricter measures to prevent unauthorized access to such data.

What would be the MOST effective solution for this issue?

Answers:

***Permission restrictions**

Data encryption

Data masking

Data obfuscation

Explanation:

Permission restrictions are an ideal option to secure data in use. Implementing permission restrictions would limit access to only authorized users, reducing the likelihood of unauthorized access.

Data encryption is an effective way to protect data at rest or in transit, but it does not typically secure data while the system processes it or when it is in active memory.

Data masking creates a structurally similar but inauthentic version of an organization's data and is not the most appropriate method for securing data in use.

Data obfuscation makes data harder to understand, making it difficult for an unauthorized user to make sense of the data. However, it is less appropriate for data in active memory compared to permission restrictions.

q_dataprot_pol_privacy_secp8

A company is updating its standard operating procedures (SOPs).

What type of policy should a company focus on to ensure compliance with data privacy regulations and outline the appropriate handling of customer data?

Answers:

***Data privacy policy**

Acceptable use policy (AUP)

Information security policy

Disaster recovery policy

Explanation:

The data privacy policy defines how the company collects, stores, processes, and shares customer data and measures taken to protect customer information. It ensures that the organization complies with relevant data protection laws and regulations, safeguarding sensitive data and promoting transparency in data handling practices.

An acceptable use policy outlines appropriate computer and network usage but doesn't cover data privacy regulations or customer information management.

An information security policy is essential to ensure that all users follow the guidelines related to information security. However, it may not cover data privacy regulations and customer data management.

A disaster recovery policy primarily deals with the steps to recover from catastrophic events such as natural disasters or security breaches.

q_dataprot_pol_retention_secp8

In a high-security environment, which of the following is the MOST important concern when removable media is no longer needed?

Answers:

Labeling

***Destruction**

Reuse

Purging

Explanation:

The most important concern is the destruction of the media. In a high-security environment, removable media is not reused. After the media is no longer needed, it must be destroyed.

Labeling is important, but it is important before removable media is put into use, not after.

Reuse and purging are not secure activities in a high-security environment. Reusing media can result in confidentiality compromise. Purging is rarely sufficient to fully remove data.

q_dataprot_pol_sovereignty_01_secp8

Which of the following is a correct interpretation of data sovereignty?

Answers:

***A jurisdiction can restrict or prevent processing and storage of data on systems that do not physically reside within that jurisdiction.**

The physical location of a data center has no bearing on the jurisdiction and laws applicable to the data stored within it.

Data can freely move across borders with no restrictions or regulations.

All data is inherently owned by the organization or individual who created it, regardless of where it is stored or processed.

Explanation:

Data sovereignty is the principle that a jurisdiction can impose restrictions or prevent the processing and storage of data on systems that do not physically reside within that jurisdiction. It often requires organizations to use location-specific storage facilities or cloud services.

The physical location of a data center often directly determines the jurisdiction and laws applicable to the data stored within it.

Data sovereignty is the opposite of this, often placing restrictions and regulations on how and where the user can move and store data.

While data ownership is a complex issue and can depend on various factors, data sovereignty refers to the laws and regulations on how and where the user can move and store data.

q_dataprot_pol_sovereignty_02_secp8

Planning to store data from various global branches, an international company is assessing the legal and regulatory compliance requirements for data storage and usage.

What should the organization consider in its analysis of government requirements?

Answers:

***Data sovereignty**

Data transfer rules

Local privacy laws

Geographic restrictions

Explanation:

Data sovereignty is the concept that data is subject to the laws and governance structures within the nation the company collects data. Complying with data sovereignty laws is crucial for an international company storing data from different locations.

Data transfer rules refer to the regulations governing the movement of data. The primary concern for an international company would be the specific laws in each country where it stores data.

Local privacy laws are a part of the considerations under data sovereignty. However, data sovereignty encompasses more than just local privacy laws.

Geographic restrictions refer to data management limitations based on geographic location. The term does not encapsulate the full range of legal and regulatory considerations.

A.0 CompTIA Security+ SY0-701 - Practice Exams

A.1 Prepare for CompTIA Security+ SY0-701 Certification

It is important to prepare for an exam by studying course material, practicing skills, and committing new concepts to memory. You can use the instructions and tests in this course to help you prepare more efficiently.

We recommend that you take the following steps as you prepare for the CompTIA Security+ SY0-701 exam:

Step	Description
Step 1: Study the course material	**Study the course materials for each section** . The course materials include text lessons, demonstrations, video lessons, and hands-on labs. As you work through the course, follow these hints for the most effective study: Review the learning and exam objectives on each section page. The objectives outline the knowledge and skills you will need for the official certification exam. Watch the videos. Watch the demonstrations. Read all text lesson fact pages. Practice the tasks in the lab simulations until you feel comfortable with your ability to complete them. Even if you already know the material, a review will be helpful when preparing for an exam.
Step 2: Take the section practice quizzes	**Pass all the Section Quizzes and be able to explain why the answers are correct.** The Section Quizzes at the end of each section will help you assess your understanding of the content for that section. Use the immediate feedback to go back and study the course material covering the questions you missed. After you have mastered the material in a section, move on to the next section.
Step 3: Study the domain review questions	**Study the Domain Review questions for each domain.** When you have finished studying the course material and passed all the Section Quizzes, you are ready to focus on exam preparation and review questions by certification exam domain. There are two types available:

	20 Questions - Twenty questions are randomly selected from the available pool of questions for a specific exam domain. All Questions - All questions from the available pool of questions for a specific exam domain are presented. This option allows the review of all available questions by exam domain.
Step 4: Take the certification practice exam	**Pass the certification practice exam with at least a 95% passing score.** After you are confident with your ability to answer each question, take the certification practice exam to assess your preparedness to take the certification exam. This exam has roughly the same number of questions and time limit as the certification exam. Check your answers after each exam, then review the course material for questions that you missed. Practice questions are designed to assess your knowledge as it relates to the exam objectives. Focus your time on understanding the topics covered in the objectives and not on memorizing answers, as the actual certification exam will have a different set of questions. It is recommended that you pass the certification practice exam multiple times as you will receive random questions each time. When you feel confident in your understanding of the exam objectives and topics, the next step is to take the certification exam.
Step 5: Schedule and take the certification exam	When you have completed the previous steps and feel prepared, schedule the CompTIA Security+ SY0-701 exam through Pearson VUE. Details on how to schedule an exam are provided in this section.

A.1.1 Security+ SY0-701 Exam Objectives

The Security Pro course and certification exam cover the following CompTIA Security+ SY0-701 objectives:

#	Domain	Module.Section
1.0	**General Security Concepts**	
1.1	Compare and contrast various types of security controls 1.1.1 - Categories 1.1.1.1 - Technical	1.1, 1.2 5.10 7.1

	1.1.1.2 - Managerial 1.1.1.3 - Operational 1.1.1.4 - Physical 1.1.2 - Control types 1.1.2.1 - Preventive 1.1.2.2 - Deterrent 1.1.2.3 - Detective 1.1.2.4 - Corrective 1.1.2.5 - Compensating 1.1.2.6 - Directive	
1.2	Summarize fundamental security concepts 1.2.1 - Confidentiality, Integrity, and Availability (CIA) 1.2.2 - Non-repudiation 1.2.3 - Authentication, Authorization, and Accounting (AAA) 1.2.3.1 - Authenticating people 1.2.3.2 - Authenticating systems 1.2.3.3 - Authorization models 1.2.4 - Gap analysis 1.2.5 - Zero trust 1.2.5.1 - Control plane 1.2.5.1.1 - Adaptive identity 1.2.5.1.2 - Threat scope reduction 1.2.5.1.3 - Policy-driven access control 1.2.5.1.4 - Policy Administrator 1.2.5.1.5 - Policy Engine	1.1 4.1, 4.2, 4.3, 4.4, 4.5 5.2 6.1 8.2 10.4

	1.2.5.2 - Data plane 1.2.5.2.1 - Implicit trust zones 1.2.5.2.2 - Subject/System 1.2.5.2.3 - Policy enforcement point 1.2.6 - Physical security 1.2.6.1 - Bollards 1.2.6.2 - Access control vestibule 1.2.6.3 - Fencing 1.2.6.4 - Video surveillance 1.2.6.5 - Security guard 1.2.6.6 - Access badge 1.2.6.7 - Lighting 1.2.6.8 - Sensors 1.2.6.8.1 - Infrared 1.2.6.8.2 - Pressure 1.2.6.8.3 - Microwave 1.2.6.8.4 - Ultrasonic 1.2.7 - Deception and disruption technology 1.2.7.1 - Honeypot 1.2.7.2 - Honeynet 1.2.7.3 - Honeyfile 1.2.7.4 - Honeytoken	
1.3	Explain the importance of change management processes and the impact to security 1.3.1 - Business processes impacting security operation 1.3.1.1 - Approval process	8.9 11.2

	1.3.1.2 - Ownership 1.3.1.3 - Stakeholders 1.3.1.4 - Impact analysis 1.3.1.5 - Test results 1.3.1.6 - Backout plan 1.3.1.7 - Maintenance window 1.3.1.8 - Standard operating procedure 1.3.2 - Technical implications 1.3.2.1 - Allow lists/deny lists 1.3.2.2 - Restricted activities 1.3.2.3 - Downtime 1.3.2.4 - Service restart 1.3.2.5 - Application restart 1.3.2.6 - Legacy applications 1.3.2.7 - Dependencies 1.3.3 - Documentation 1.3.3.1 - Updating diagrams 1.3.3.2 - Updating policies/procedures 1.3.4 -Version control	
1.4	Explain the importance of using appropriate cryptographic solutions 1.4.1 - Public key infrastructure (PKI) 1.4.1.1 - Public key 1.4.1.2 - Private key 1.4.1.3 - Key escrow 1.4.2 - Encryption	3.1, 3.2, 3.3, 3.4, 3.5 4.3 5.8 6.6 7.3 8.1, 8.2, 8.7 10.8

- 1.4.2.1 - Level
  - 1.4.2.1.1 - Full-disk
  - 1.4.2.1.2 - Partition
  - 1.4.2.1.3 - File
  - 1.4.2.1.4 - Volume
  - 1.4.2.1.5 - Database
  - 1.4.2.1.6 - Record
- 1.4.2.2 - Transport/communication
- 1.4.2.3 - Asymmetric
- 1.4.2.4 - Symmetric
- 1.4.2.5 - Key exchange
- 1.4.2.6 - Algorithms
- 1.4.2.7 - Key length

1.4.3 - Tools

- 1.4.3.1 - Trusted Platform Module (TPM)
- 1.4.3.2 - Hardware security module (HSM)
- 1.4.3.3 - Key management system
- 1.4.3.4 - Secure enclave

1.4.4 - Obfuscation

- 1.4.4.1 - Steganography
- 1.4.4.2 - Tokenization
- 1.4.4.3 - Data masking

1.4.5 - Hashing

1.4.6 - Salting

1.4.7 - Digital signatures

	1.4.8 - Key stretching 1.4.9 - Blockchain 1.4.10 - Open public ledger 1.4.11 - Certificates 1.4.11.1 - Certificate authorities 1.4.11.2 - Certificate revocation lists (CRLs) 1.4.11.3 - Online Certificate Status Protocol (OCSP) 1.4.11.4 - Self-signed 1.4.11.5 - Third-party 1.4.11.6 - Root of trust 1.4.11.7 - Certificate signing request (CSR) generation 1.4.11.8 - Wildcard	
2.0	**Threats, Vulnerabilities, and Mitigations**	
2.1	Compare and contrast common threat actors and motivations 2.1.1 - Threat actors 2.1.1.1 - Nation-state 2.1.1.2 - Unskilled attacker 2.1.1.3 - Hacktivist 2.1.1.4 - Insider threat 2.1.1.5 - Organized crime 2.1.1.6 - Shadow IT 2.1.2 - Attributes of actors 2.1.2.1 - Internal/external 2.1.2.2 - Resources/funding	1.1 2.1, 2.2 6.5

	2.1.2.3 - Level of sophistication/capability 2.1.3 - Motivations 2.1.3.1 - Data exfiltration 2.1.3.2 - Espionage 2.1.3.3 - Service disruption 2.1.3.4 - Blackmail 2.1.3.5 - Financial gain 2.1.3.6 - Philosophical/political beliefs 2.1.3.7- Ethical 2.1.3.8 - Revenge 2.1.3.9 - Disruption/chaos 2.1.3.10 - War	
2.2	Explain common threat vectors and attack surfaces 2.2.1 - Message-based 2.2.1.1 - Email 2.2.1.2 - Short Message Service (SMS) 2.2.1.3 - Instant messaging (IM) 2.2.2 - Image-based 2.2.3 - File-based 2.2.4 - Voice call 2.2.5 - Removable device 2.2.6 - Vulnerable software 2.2.6.1 - Client-based vs. agentless 2.2.7 - Unsupported systems and applications 2.2.8 - Unsecure networks	2.1, 2.2 5.10 6.2, 6.6 7.2, 7.4 8.1, 8.2, 8.3, 8.5, 8.6, 8.9 10.4, 10.7, 10.8

	2.2.8.1 - Wireless 2.2.8.2 - Wired 2.2.8.3 - Bluetooth 2.2.9 - Open service ports 2.2.10 - Default credentials 2.2.11 - Supply chain 2.2.11.1 - Managed service providers (MSPs) 2.2.11.2 - Vendors 2.2.11.3 - Suppliers 2.2.12 - Human vectors/social engineering 2.2.12.1 - Phishing 2.2.12.2 - Vishing 2.2.12.3 - Smishing 2.2.12.4 - Misinformation/disinformation 2.2.12.5 - Impersonation 2.2.12.6 - Business email compromise 2.2.12.7 - Pretexting 2.2.12.8 - Watering hole 2.2.12.9 - Brand impersonation 2.2.12.10 - Typosquatting	
2.3	Explain various types of vulnerabilities 2.3.1 - Application 2.3.1.1 - Memory injection 2.3.1.2 - Buffer overflow 2.3.1.3 - Race conditions	2.1 5.7, 5.10 7.1 8.6, 8.8 10.1, 10.4, 10.5

- 2.3.1.3.1 - Time-of-check (TOC)
- 2.3.1.3.2 - Time-of-use (TOU)
- 2.3.1.4 - Malicious update

2.3.2 - Operating system (OS)-based

2.3.3 - Web-based

- 2.3.3.1 - Structured Query Language injection (SQLi)
- 2.3.3.2 - Cross-site scripting (XSS)

2.3.4 - Hardware

- 2.3.4.1 - Firmware
- 2.3.4.2 - End-of-life
- 2.3.4.3 - Legacy

2.3.5 - Virtualization

- 2.3.5.1 - Virtual machine (VM) escape
- 2.3.5.2 - Resource reuse

2.3.6 - Cloud-specific

2.3.7 - Supply chain

- 2.3.7.1 - Service provider
- 2.3.7.2 - Hardware provider
- 2.3.7.3 - Software provider

2.3.8 - Cryptographic

2.3.9 - Misconfiguration

2.3.10 - Mobile device

- 2.3.10.1 - Side loading
- 2.3.10.2 - Jailbreaking

2.3.11 - Zero-day

2.4	Given a scenario, analyze indicators of malicious activity 2.4.1 - Malware attacks 2.4.1.1 - Ransomware 2.4.1.2 - Trojan 2.4.1.3 - Worm 2.4.1.4 - Spyware 2.4.1.5 - Bloatware 2.4.1.6 - Virus 2.4.1.7 - Keylogger 2.4.1.8 - Logic bomb 2.4.1.9 - Rootkit 2.4.2 - Physical attacks 2.4.2.1 - Brute force 2.4.2.2 - Radio frequency identification (RFID) cloning 2.4.2.3 - Environmental 2.4.3 - Network attacks 2.4.3.1 - Distributed denial-of-service (DDoS) 2.4.3.1.1 - Amplified 2.4.3.1.2 - Reflected 2.4.3.2 - Domain Name System (DNS) attacks 2.4.3.3 - Wireless 2.4.3.4 - On-path 2.4.3.5 - Credential replay 2.4.3.6 - Malicious code	2.1, 2.3 3.1 4.5 5.7 6.1, 6.4, 6.5, 6.6 8.5, 8.7, 8.8 10.9

	2.4.4 - Application attacks 2.4.4.1 - Injection 2.4.4.2 - Buffer overflow 2.4.4.3 - Replay 2.4.4.4 - Privilege escalation 2.4.4.5 - Forgery 2.4.4.6 - Directory traversal 2.4.5 - Cryptographic attacks 2.4.5.1 - Downgrade 2.4.5.2 - Collision 2.4.5.3 - Birthday 2.4.6 - Password attacks 2.4.6.1 - Spraying 2.4.6.2 - Brute force 2.4.7 - Indicators 2.4.7.1 - Account lockout 2.4.7.2 - Concurrent session usage 2.4.7.3 - Blocked content 2.4.7.4 - Impossible travel 2.4.7.5 - Resource consumption 2.4.7.6 - Resource inaccessibility 2.4.7.7 - Out-of-cycle logging 2.4.7.8 - Published/documented 2.4.7.9 - Missing logs	
2.5	Explain the purpose of mitigation techniques used to secure the enterprise	2.1 3.4 4.1, 4.3, 4.6

	2.5.1 - Segmentation 2.5.2 - Access control 2.5.2.1 - Access control list (ACL) 2.5.2.2 - Permissions 2.5.3 - Application allow list 2.5.4 - Isolation 2.5.5 - Patching 2.5.6 - Encryption 2.5.7 - Monitoring 2.5.8 - Least privilege 2.5.9 - Configuration enforcement 2.5.10 - Decommissioning 2.5.11 - Hardening techniques 2.5.11.1 - Encryption 2.5.11.2 - Installation of endpoint protection 2.5.11.3 - Host-based firewall 2.5.11.4 - Host-based intrusion prevention system (HIPS) 2.5.11.5 - Disabling ports/protocols 2.5.11.6 - Default password changes 2.5.11.7 - Removal of unnecessary software	5.4, 5.7, 5.10 6.2, 6.3, 6.4 7.3 8.1, 8.2, 8.3, 8.6, 8.8, 8.9 9.1, 9.2 10.1, 10.5, 10.6, 10.7 13.2
3.0	**Security Architecture**	
3.1	Compare and contrast security implications of different architecture models 3.1.1 - Architecture and infrastructure concepts 3.1.1.1 - Cloud	5.1 9.4 10.1, 10.2, 10.3, 10.4, 10.6, 10.8

- 3.1.1.1.1 - Responsibility matrix
- 3.1.1.1.2 - Hybrid considerations
- 3.1.1.1.3 - Third-party vendors
- 3.1.1.2 - Infrastructure as code (IaC)
- 3.1.1.3 - Serverless
- 3.1.1.4 - Microservices
- 3.1.1.5 - Network infrastructure
- 3.1.1.5.1 - Physical isolation
- 3.1.1.5.1.1 - Air-gapped
- 3.1.1.5.2 - Logical segmentation
- 3.1.1.5.3 - Software-defined networking (SDN)
- 3.1.1.6 - On-premises
- 3.1.1.7 - Centralized/decentralized
- 3.1.1.8 - Containerization
- 3.1.1.9 - Virtualization
- 3.1.1.10 - IoT
- 3.1.1.11 - Industrial control systems (ICS)/supervisory control and data acquisition (SCADA)
- 3.1.1.12 - Real-time operating system (RTOS)
- 3.1.1.13 - Embedded systems
- 3.1.1.14 - High availability

3.1.2 - Considerations

- 3.1.2.1 - Availability
- 3.1.2.2 - Resilience
- 3.1.2.3 - Cost

	3.1.2.4 - Responsiveness 3.1.2.5 - Scalability 3.1.2.6 - Ease of deployment 3.1.2.7 - Risk transference 3.1.2.8 - Ease of recovery 3.1.2.9 - Patch availability 3.1.2.10 - Inability to patch 3.1.2.11 - Power 3.1.2.12 - Compute	
3.2	Given a scenario, apply security principles to secure enterprise infrastructure 3.2.1 - Infrastructure considerations 3.2.1.1 - Device placement 3.2.1.2 - Security zones 3.2.1.3 - Attack surface 3.2.1.4 - Connectivity 3.2.1.5 - Failure modes 3.2.1.5.1 - Fail-open 3.2.1.5.2 - Fail-closed 3.2.1.6 - Device attribute 3.2.1.6.1 - Active vs. passive 3.2.1.6.2 - Inline vs. tap/monitor 3.2.1.7 - Network appliances 3.2.1.7.1 - Jump server 3.2.1.7.2 - Proxy server 3.2.1.7.3 - Intrusion prevention system (IPS)/intrusion detection system (IDS)	2.1 4.8 5.1, 5.2, 5.3, 5.4, 5.5, 5.9, 5.10 6.3 7.3 8.2, 8.6, 8.7, 8.8 9.2, 9.4 10.1, 10.2, 10.9

	3.2.1.7.4 - Load balancer 3.2.1.7.5 - Sensors 3.2.1.8 - Port security 3.2.1.8.1 - 802.1X 3.2.1.8.2 - Extensible Authentication Protocol (EAP) 3.2.1.9 - Firewall types 3.2.1.9.1 - Web application firewall (WAF) 3.2.1.9.2 - Unified threat management (UTM) 3.2.1.9.3 - Next-generation firewall (NGFW) 3.2.1.9.4 - Layer 4/Layer 7 3.2.2 - Secure communication/access 3.2.2.1 - Virtual private network (VPN) 3.2.2.2 - Remote access 3.2.2.3 - Tunneling 3.2.2.3.1 - Transport Layer Security (TLS) 3.2.2.3.2 - Internet protocol security (IPSec) 3.2.2.4 - Software-defined wide area network (SD-WAN) 3.2.2.5 - Secure access service edge (SASE) 3.2.3 - Selection of effective controls	
3.3	Compare and contrast concepts and strategies to protect data 3.3.1 - Data types 3.3.1.1 - Regulated 3.3.1.2 - Trade secret 3.3.1.3 - Intellectual property	5.9, 5.10 7.3 8.9 10.5, 10.8 13.1, 13.2

- 3.3.1.4 - Legal information
- 3.3.1.5 - Financial information
- 3.3.1.6 - Human and non-human readable

3.3.2 - Data classifications

- 3.3.2.1 - Sensitive
- 3.3.2.2 - Confidential
- 3.3.2.3 - Public
- 3.3.2.4 - Restricted
- 3.3.2.5 - Private
- 3.3.2.6 - Critical

3.3.3 - General data considerations

- 3.3.3.1 - Data states
  - 3.3.3.1.1 - Data at rest
  - 3.3.3.1.2 - Data in transit
  - 3.3.3.1.3 - Data in use
- 3.3.3.2 - Data sovereignty
- 3.3.3.3 - Geolocation

3.3.4 - Methods to secure data

- 3.3.4.1 - Geographic restrictions
- 3.3.4.2 - Encryption
- 3.3.4.3 - Hashing
- 3.3.4.4 - Masking
- 3.3.4.5 - Tokenization
- 3.3.4.6 - Obfuscation
- 3.3.4.7 - Segmentation

	3.3.4.8 - Permission restrictions	
3.4	Explain the importance of resilience and recovery in security architecture 3.4.1 - High availability 3.4.1.1 - Load balancing vs. clustering 3.4.2 - Site considerations 3.4.2.1 - Hot 3.4.2.2 - Cold 3.4.2.3 - Warm 3.4.2.4 - Geographic dispersion 3.4.3 - Platform diversity 3.4.4 - Multi-cloud systems 3.4.5 - Continuity of operations 3.4.6 - Capacity planning 3.4.6.1 - People 3.4.6.2 - Technology 3.4.6.3 - Infrastructure 3.4.7 - Testing 3.4.7.1 - Tabletop exercises 3.4.7.2 - Fail over 3.4.7.3 - Simulation 3.4.7.4 - Parallel processing 3.4.8 - Backups 3.4.8.1 - Onsite/offsite 3.4.8.2 - Frequency 3.4.8.3 - Encryption	6.1 9.4, 9.5 10.1, 10.4 12.1

	3.4.8.4 - Snapshots 3.4.8.5 - Recovery 3.4.8.6 - Replication 3.4.8.7 - Journaling 3.4.9 - Power 3.4.9.1 - Generators 3.4.9.2 - Uninterruptible power supply (UPS)	
4.0	**Security Operations**	
4.1	Given a scenario, apply common security techniques to computing resources 4.1.1 - Secure baselines 4.1.1.1 - Establish 4.1.1.2 - Deploy 4.1.1.3 - Maintain 4.1.2 - Hardening targets 4.1.2.1 - Mobile devices 4.1.2.2 - Workstations 4.1.2.3 - Switches 4.1.2.4 - Routers 4.1.2.5 - Cloud infrastructure 4.1.2.6 - Servers 4.1.2.7 - ICS/SCADA 4.1.2.8 - Embedded systems 4.1.2.9 - RTOS 4.1.2.10 - IoT devices 4.1.3. Wireless devices	3.5 4.6, 4.8 5.9, 5.10 6.2, 6.4 7.3 8.1, 8.2, 8.4, 8.5, 8.6, 8.8, 8.9 9.2 10.1, 10.2, 10.4, 10.5, 10.6, 10.7, 10.8, 10.9 13.2

4.1.3.1 - Installation considerations

4.1.3.1.1 - Site surveys

4.1.3.1.2 - Heat maps

4.1.4 - Mobile solutions

4.1.4.1 - Mobile device management (MDM)

4.1.4.2 - Deployment models

4.1.4.2.1 - Bring your own device (BYOD)

4.1.4.2.2 - Corporate-owned, personally enabled (COPE)

4.1.4.2.3 - Choose your own device (CYOD)

4.1.4.3 - Connections methods

4.1.4.3.1 - Cellular

4.1.4.3.2 - Wi-Fi

4.1.4.3.3 - Bluetooth

4.1.5 - Wireless security settings

4.1.5.1 - Wi-Fi Protected Access 3 (WPA3)

4.1.5.2 - AAA/Remote Authentication Dial-In User Service (RADIUS)

4.1.5.3 - Cryptographic protocols

4.1.5.4 - Authentication protocols

4.1.6 - Application security

4.1.6.1 - Input validation

4.1.6.2 - Secure cookies

4.1.6.3 - Static code analysis

4.1.6.4 - Code signing

4.1.7. Sandboxing

4.1.8. Monitoring

4.2	Explain the security implications of proper hardware, software, and data asset management 4.2.1 - Acquisition/procurement process 4.2.2 - Assignment/accounting 4.2.2.1 - Ownership 4.2.2.2 - Classification 4.2.3 - Disposal/decommissioning 4.2.3.1 - Sanitization 4.2.3.2 - Destruction 4.2.3.3 - Certification 4.2.3.4 - Data retention	8.2 10.6 12.1 13.1, 13.2
4.3	Explain various activities associated with vulnerability management 4.3.1 - Identification methods 4.3.1.1 - Vulnerability scan 4.3.1.2 - Application security 4.3.1.2.1 - Static analysis 4.3.1.2.2 - Dynamic analysis 4.3.1.2.3 - Package monitoring 4.3.1.3 - Threat feed 4.3.1.3.1 - Open-source intelligence (OSINT) 4.3.1.3.2 - Proprietary/third-party 4.3.1.3.3 - Information-sharing organization 4.3.1.3.4 - Dark web 4.3.1.4 - Penetration testing 4.3.1.5 - Responsible disclosure program 4.3.1.5.1 - Bug bounty program	5.8 6.2, 6.3 7.1, 7.2, 7.3, 7.4 8.1, 8.2, 8.9 9.2 10.1 12.3

	4.3.1.6 - System/process audit 4.3.2 - Analysis 4.3.2.1 - Confirmation 4.3.2.1.1 - False positive 4.3.2.1.2 - False negative 4.3.2.2 - Prioritize 4.3.2.3 - Common Vulnerability Scoring System (CVSS) 4.3.2.4 - Common Vulnerability Enumeration (CVE) 4.3.2.5 - Vulnerability classification 4.3.2.6 - Exposure factor 4.3.2.7 - Environmental variables 4.3.2.8 - Industry/organizational impact 4.3.2.9 - Risk tolerance 4.3.3 - Vulnerability response and remediation 4.3.3.1 - Patching 4.3.3.2 - Insurance 4.3.3.3 - Segmentation 4.3.3.4 - Compensating controls 4.3.3.5 - Exceptions and exemptions 4.3.4 - Validation of remediation 4.3.4.1 - Rescanning 4.3.4.2 - Audit 4.3.4.3 - Verification 4.3.5 - Reporting	
4.4	Explain security alerting and monitoring concepts and tools	6.2

	4.4.1 - Monitoring computing resources 4.4.1.1 - Systems 4.4.1.2 - Applications 4.4.1.3 - Infrastructure 4.4.2 - Activities 4.4.2.1 - Log aggregation 4.4.2.2 - Alerting 4.4.2.3 - Scanning 4.4.2.4 - Reporting 4.4.2.5 - Archiving 4.4.2.6 - Alert response and remediation/validation 4.4.2.6.1 - Quarantine 4.4.2.6.2 - Alert tuning 4.4.3 - Tools 4.4.3.1 - Security Content Automation Protocol (SCAP) 4.4.3.2 - Benchmarks 4.4.3.3 - Agents/agentless 4.4.3.3.1 - Security information and event management (SIEM) 4.4.3.3.2 - Antivirus 4.4.3.3.3 - Data loss prevention (DLP) 4.4.3.4 - Simple Network Management Protocol (SNMP) traps 4.4.3.5 - NetFlow 4.4.3.6 - Vulnerability scanners	7.1, 7.2, 7.3, 7.4 9.2 12.3 13.2

4.5	Given a scenario, modify enterprise capabilities to enhance security 4.5.1 - Firewall 4.5.1.1 - Rules 4.5.1.2 - Access lists 4.5.1.3 - Ports/protocols 4.5.1.4 - Screened subnets 4.5.2 - IDS/IPS 4.5.2.1 - Trends 4.5.2.2 - Signatures 4.5.3 - Web filter 4.5.3.1 - Agent-based 4.5.3.2 - Centralized proxy 4.5.3.3 - Universal Resource Locator (URL) scanning 4.5.3.4 - Content categorization 4.5.3.5 - Block rules 4.5.3.6 - Reputation 4.5.4 - Operating system security 4.5.4.1 - Group Policy 4.5.4.2 - SELinux 4.5.5 - Implementation of secure protocols 4.5.5.1 - Protocol selection 4.5.5.2 - Port selection 4.5.5.3 - Transport method 4.5.6 - DNS filtering	4.4, 4.5, 4.6 5.2, 5.3, 5.4, 5.6, 5.9, 5.10 6.2, 6.3, 6.4 8.1, 8.2, 8.9 10.5, 10.7, 10.9 12.3

	4.5.7 - Email security 4.5.7.1 - Domain-based Message Authentication Reporting and Conformance (DMARC) 4.5.7.2 - DomainKeys Identified Mail (DKIM) 4.5.7.3 - Sender Policy Framework (SPF) 4.5.7.4 - Gateway 4.5.8. File integrity monitoring 4.5.9. DLP 4.5.10. Network access control (NAC) 4.5.11. Endpoint detection and response (EDR)/extended detection and response (XDR) 4.5.12. User behavior analytics	
4.6	Given a scenario, implement and maintain identity and access management 4.6.1 - Provisioning/de-provisioning user accounts 4.6.2 - Permission assignments and implications 4.6.3 - Identity proofing 4.6.4 - Federation 4.6.5 - Single sign-on (SSO) 4.6.5.1 - Lightweight Directory Access Protocol (LDAP) 4.6.5.2 - Open authorization (OAuth) 4.6.5.3 - Security Assertions Markup Language (SAML) 4.6.6 - Interoperability 4.6.7 - Attestation 4.6.8 - Access controls 4.6.8.1 - Mandatory	4.1, 4.2, 4.3, 4.5, 4.6, 4.7, 4.9 5.7 6.1 8.1, 8.2, 8.8, 8.9 10.5, 10.7 11.1 13.2

4.6.8.2 - Discretionary

4.6.8.3 - Role-based

4.6.8.4 - Rule-based

4.6.8.5 - Attribute-based

4.6.8.6 - Time-of-day restrictions

4.6.8.7 - Least privilege

4.6.9 - Multifactor authentication

4.6.9.1 - Implementations

4.6.9.1.1 - Biometrics

4.6.9.1.2 - Hard/soft authentication tokens

4.6.9.1.3 - Security keys

4.6.9.2 - Factors

4.6.9.2.1 - Something you know

4.6.9.2.2 - Something you have

4.6.9.2.3 - Something you are

4.6.9.2.4 - Somewhere you are

4.6.10 - Password concepts

4.6.10.1 - Password best practices

4.6.10.1.1 - Length

4.6.10.1.2 - Complexity

4.6.10.1.3 - Reuse

4.6.10.1.4 - Expiration

4.6.10.1.5 - Age

4.6.10.2 - Password managers

4.6.10.3 - Passwordless

	4.6.11 - Privileged access management tools 4.6.11.1 - Just-in-time permissions 4.6.11.2 - Password vaulting 4.6.11.3 - Temporal accounts	
4.7	Explain the importance of automation and orchestration related to secure operations 4.7.1 - Use cases of automation and scripting 4.7.1.1 - User provisioning 4.7.1.2 - Resource provisioning 4.7.1.3 - Guard rails 4.7.1.4 - Security groups 4.7.1.5 - Ticket creation 4.7.1.6 - Escalation 4.7.1.7 - Enabling/disabling services and access 4.7.1.8 - Continuous integration and testing 4.7.1.9 - Integrations and Application programming interfaces (APIs) 4.7.2 - Benefits 4.7.2.1 - Efficiency/time saving 4.7.2.2 - Enforcing baselines 4.7.2.3 - Standard infrastructure configurations 4.7.2.4 - Scaling in a secure manner 4.7.2.5 - Staff retention 4.7.2.6 - Reaction time 4.7.2.7 - Workforce multiplier 4.7.3 - Other considerations	6.5 8.1, 8.9 11.3

	4.7.3.1 - Complexity 4.7.3.2 - Cost 4.7.3.3 - Single point of failure 4.7.3.4 - Technical debt 4.7.3.5 - Ongoing supportability	
4.8	Explain appropriate incident response activities 4.8.1 - Process 4.8.1.1 - Preparation 4.8.1.2 - Detection 4.8.1.3 - Analysis 4.8.1.4 - Containment 4.8.1.5 - Eradication 4.8.1.6 - Recovery 4.8.1.7 - Lessons learned 4.8.2 - Training 4.8.3 - Testing 4.8.3.1 - Tabletop exercise 4.8.3.2 - Simulation 4.8.4 - Root cause analysis 4.8.5 - Threat hunting 4.8.6 - Digital forensics 4.8.6.1 - Legal hold 4.8.6.2 - Chain of custody 4.8.6.3 - Acquisition 4.8.6.4 - Reporting	7.1 9.1, 9.3 13.2

	4.8.6.5 - Preservation 4.8.6.6 - E-discovery	
4.9	Given a scenario, use data sources to support an investigation 4.9.1 - Log data 4.9.1.1 - Firewall logs 4.9.1.2 - Application logs 4.9.1.3 - Endpoint logs 4.9.1.4 - OS-specific security logs 4.9.1.5 - IPS/IDS logs 4.9.1.6 - Network logs 4.9.1.7 - Metadata 4.9.2 - Data sources 4.9.2.1 - Vulnerability scans 4.9.2.2 - Automated reports 4.9.2.3 - Dashboards 4.9.2.4 - Packet captures	6.2, 6.4, 6.5 7.1, 7.2, 7.3 9.2, 9.3 12.3
5.0	**Security Program Management and Oversight**	
5.1	Summarize elements of effective security governance 5.1.1 - Guidelines 5.1.2 - Policies 5.1.2.1 - Acceptable use policy (AUP) 5.1.2.2 - Information security policies 5.1.2.3 - Business continuity 5.1.2.4 - Disaster recovery 5.1.2.5 - Incident response 5.1.2.6 - Software development lifecycle (SDLC)	4.1 5.7, 5.10 7.3 8.9 9.1 10.7 11.1, 11.2 12.1, 12.2, 12.3 13.1, 13.2

- 5.1.2.7 - Change management
- 5.1.3 - Standards
  - 5.1.3.1 - Password
  - 5.1.3.2 - Access control
  - 5.1.3.3 - Physical security
  - 5.1.3.4 - Encryption
- 5.1.4 - Procedures
  - 5.1.4.1 - Change management
  - 5.1.4.2 - Onboarding/offboarding
  - 5.1.4.3 - Playbooks
- 5.1.5 - External considerations
  - 5.1.5.1 - Regulatory
  - 5.1.5.2 - Legal
  - 5.1.5.3 - Industry
  - 5.1.5.4 - Local/regional
  - 5.1.5.5 - National
  - 5.1.5.6 - Global
- 5.1.6 - Monitoring and revision
- 5.1.7 - Types of governance structures
  - 5.1.7.1 - Boards
  - 5.1.7.2 - Committees
  - 5.1.7.3 - Government entities
  - 5.1.7.4 - Centralized/decentralized
- 5.1.8 - Roles and responsibilities for systems and data
  - 5.1.8.1 - Owners

	5.1.8.2 - Controllers 5.1.8.3 - Processors 5.1.8.4 - Custodians/stewards	
5.2	Explain elements of the risk management process 5.2.1 - Risk identification 5.2.2 - Risk assessment 5.2.2.1 - Ad hoc 5.2.2.2 - Recurring 5.2.2.3 - One-time 5.2.2.4 - Continuous 5.2.3 - Risk analysis 5.2.3.1 - Qualitative 5.2.3.2 - Quantitative 5.2.3.3 - Single loss expectancy (SLE) 5.2.3.4 - Annualized loss expectancy (ALE) 5.2.3.5 - Annualized rate of occurrence (ARO) 5.2.3.6 - Probability 5.2.3.7 - Likelihood 5.2.3.8 - Exposure factor 5.2.3.9 - Impact 5.2.4 - Risk register 5.2.4.1 - Key risk indicators 5.2.4.2 - Risk owners 5.2.4.3 - Risk threshold 5.2.5 - Risk tolerance	8.9 11.2 12.1

	5.2.6 - Risk appetite 5.2.6.1 - Expansionary 5.2.6.2 - Conservative 5.2.6.3 - Neutral 5.2.7 - Risk management strategies 5.2.7.1 - Transfer 5.2.7.2 - Accept 5.2.7.2.1 - Exemption 5.2.7.2.2 - Exception 5.2.7.3 - Avoid 5.2.7.4 - Mitigate 5.2.8 - Risk reporting 5.2.9 - Business impact analysis 5.2.9.1 - Recovery time objective (RTO) 5.2.9.2 - Recovery point objective (RPO) 5.2.9.3 - Mean time to repair (MTTR) 5.2.9.4 - Mean time between failures (MTBF)	
5.3	Explain the processes associated with third-party risk assessment and management 5.3.1 - Vendor assessment 5.3.1.1 - Penetration testing 5.3.1.2 - Right-to-audit clause 5.3.1.3 - Evidence of internal audits 5.3.1.4 - Independent assessments 5.3.1.5 - Supply chain analysis 5.3.2 - Vendor selection	7.4 9.4 10.4 12.2 13.2

	5.3.2.1 - Due diligence 5.3.2.2 - Conflict of interest 5.3.3 - Agreement types 5.3.3.1 - Service-level agreement (SLA) 5.3.3.2 - Memorandum of agreement (MOA) 5.3.3.3 - Memorandum of understanding (MOU) 5.3.3.4 - Master service agreement (MSA) 5.3.3.5 - Work order (WO)/statement of work (SOW) 5.3.3.6 - Non-disclosure agreement (NDA) 5.3.3.7 - Business partners agreement (BPA) 5.3.4 - Vendor monitoring 5.3.5 - Questionnaires 5.3.6 - Rules of engagement	
5.4	Summarize elements of effective security compliance 5.4.1 - Compliance reporting 5.4.1.1 - Internal 5.4.1.2 - External 5.4.2 - Consequences of non-compliance 5.4.2.1 - Fines 5.4.2.2 - Sanctions 5.4.2.3 - Reputational damage 5.4.2.4 - Loss of license 5.4.2.5 - Contractual impacts 5.4.3 - Compliance monitoring 5.4.3.1 - Due diligence/care	6.2 13.1, 13.2

	5.4.3.2 - Attestation and acknowledgement 5.4.3.3 - Internal and external 5.4.3.4 - Automation 5.4.4 - Privacy 5.4.4.1 - Legal implications 5.4.4.1.1 - Local/regional 5.4.4.1.2 - National 5.4.4.1.3 - Global 5.4.4.2 - Data subject 5.4.4.3 - Controller vs. processor 5.4.4.4 - Ownership 5.4.4.5 - Data inventory and retention 5.4.4.6 - Right to be forgotten	
5.5	Explain types and purposes of audits and assessments 5.5.1 - Attestation 5.5.2 - Internal 5.5.2.1 - Compliance 5.5.2.2 - Audit committee 5.5.2.3 - Self-assessments 5.5.3 - External 5.5.3.1 - Regulatory 5.5.3.2 - Examinations 5.5.3.3 - Assessment 5.5.3.4 - Independent third-party audit 5.5.4 - Penetration testing	6.2, 6.5 7.4 12.3

	5.5.4.1 - Physical 5.5.4.2 - Offensive 5.5.4.3 - Defensive 5.5.4.4 - Integrated 5.5.4.5 - Known environment 5.5.4.6 - Partially known environment 5.5.4.7 - Unknown environment 5.5.4.8 - Reconnaissance 5.5.4.8.1 - Passive 5.5.4.8.2 - Active	
5.6	Given a scenario, implement security awareness practices 5.6.1 - Phishing 5.6.1.1 - Campaigns 5.6.1.2 - Recognizing a phishing attempt 5.6.1.3 - Responding to reported suspicious messages 5.6.2 - Anomalous behavior recognition 5.6.2.1 - Risky 5.6.2.2 - Unexpected 5.6.2.3 - Unintentional 5.6.3 - User guidance and training 5.6.3.1 - Policy/handbooks 5.6.3.2 - Situational awareness 5.6.3.3 - Insider threat 5.6.3.4 - Password management 5.6.3.5 - Removable media and cables	2.2 6.3, 6.6 7.3 10.9 13.2

	5.6.3.6 - Social engineering 5.6.3.7 - Operational security 5.6.3.8 - Hybrid/remote work environments 5.6.4 - Reporting and monitoring 5.6.4.1 - Initial 5.6.4.2 - Recurring 5.6.5 - Development 5.6.6 - Execution	

A.1.2 Security+ SY0-701 Exam Objectives by Course Section

The Security Pro course covers the following CompTIA Security+ SY0-701 exam objectives:

Section	Title	Objectives
1.0	**Security Concepts**	
1.1	Security Introduction	1.1 Compare and contrast various types of security controls • 1.1.1 - Categories 1.1.1.1 - Technical 1.1.1.2 - Managerial 1.1.1.3 - Operational 1.1.1.4 - Physical 1.2 Summarize fundamental security concepts • 1.2.1 - Confidentiality, Integrity, and Availability (CIA) 2.1 Compare and contrast common threat actors and motivations • 2.1.1 - Threat actors 2.1.1.1 - Nation-state 2.1.1.5 - Organized crime

1.2	Security Controls	1.1 Compare and contrast various types of security controls • 1.1.1 - Categories 1.1.1.1 - Technical 1.1.1.2 - Managerial 1.1.1.3 - Operational 1.1.1.4 - Physical • 1.1.2 - Control types 1.1.2.1 - Preventive 1.1.2.2 - Deterrent 1.1.2.3 - Detective 1.1.2.4 - Corrective 1.1.2.5 - Compensating 1.1.2.6 - Directive
1.3	Use the Simulator	
2.0	**Threats, Vulnerabilities, and Mitigations**	
2.1	Understanding Attacks	2.1 Compare and contrast common threat actors and motivations • 2.1.1 - Threat actors 2.1.1.1 - Nation-state 2.1.1.2 - Unskilled attacker 2.1.1.3 - Hacktivist 2.1.1.4 - Insider threat 2.1.1.5 - Organized crime 2.1.1.6 - Shadow IT • 2.1.2 - Attributes of actors 2.1.2.1 - Internal/external

2.1.2.2 - Resources/funding

2.1.2.3 - Level of sophistication/capability

- 2.1.3 - Motivations

2.1.3.1 - Data exfiltration

2.1.3.2 - Espionage

2.1.3.3 - Service disruption

2.1.3.4 - Blackmail

2.1.3.5 - Financial gain

2.1.3.6 - Philosophical/political beliefs

2.1.3.7- Ethical

2.1.3.8 - Revenge

2.1.3.9 - Disruption/chaos

2.1.3.10 - War

2.2 Explain common threat vectors and attack surfaces

- 2.2.1 - Message-based

2.2.1.1 - Email

2.2.1.2 - Short Message Service (SMS)

2.2.1.3 - Instant messaging (IM)

- 2.2.2 - Image-based
- 2.2.3 - File-based
- 2.2.5 - Removable device
- 2.2.6 - Vulnerable software

2.2.6.1 - Client-based vs. agentless

- 2.2.7 - Unsupported systems and applications
- 2.2.8 - Unsecure networks

2.2.8.1 - Wireless

2.2.8.2 - Wired

		2.2.8.3 - Bluetooth • 2.2.9 - Open service ports • 2.2.10 - Default credentials • 2.2.11 - Supply chain 2.2.11.1 - Managed service providers (MSPs) 2.2.11.2 - Vendors 2.2.11.3 - Suppliers 2.3 Explain various types of vulnerabilities • 2.3.7 - Supply chain 2.3.7.1 - Service provider 2.3.7.2 - Hardware provider 2.3.7.3 - Software provider 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.4 - Application attacks 2.4.4.4 - Privilege escalation 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.8 - Least privilege 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.3 - Attack surface
2.2	Social Engineering	2.1 Compare and contrast common threat actors and motivations • 2.1.1 - Threat actors 2.1.1.2 - Unskilled attacker 2.1.1.4 - Insider threat • 2.1.3 - Motivations

		2.2 Explain common threat vectors and attack surfaces • 2.2.12 - Human vectors/social engineering 2.2.12.1 - Phishing 2.2.12.2 - Vishing 2.2.12.3 - Smishing 2.2.12.4 - Misinformation/disinformation 2.2.12.5 - Impersonation 2.2.12.6 - Business email compromise 2.2.12.7 - Pretexting 2.2.12.8 - Watering hole 2.2.12.9 - Brand impersonation 2.2.12.10 - Typosquatting 5.6 Given a scenario, implement security awareness practices • 5.6.3 - User guidance and training 5.6.3.6 - Social engineering • 5.6.5 - Development
2.3	Malware	2.4 Given a scenario, analyze indicators of malicious activity • 2.4.1 - Malware attacks 2.4.1.1 - Ransomware 2.4.1.2 - Trojan 2.4.1.3 - Worm 2.4.1.4 - Spyware 2.4.1.5 - Bloatware 2.4.1.6 - Virus

		2.4.1.7 - Keylogger 2.4.1.8 - Logic bomb 2.4.1.9 - Rootkit
3.0	**Cryptographic Solutions**	
3.1	Cryptography	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.1 - Public key infrastructure (PKI) 1.4.1.1 - Public key 1.4.1.2 - Private key 1.4.1.3 - Key escrow • 1.4.2 - Encryption 1.4.2.3 - Asymmetric 1.4.2.4 - Symmetric 1.4.2.6 - Algorithms 1.4.2.7 - Key length • 1.4.4 - Obfuscation 1.4.4.1 - Steganography • 1.4.5 - Hashing • 1.4.6 - Salting • 1.4.7 - Digital signatures • 1.4.9 - Blockchain 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.5 - Cryptographic attacks 2.4.5.1 - Downgrade 2.4.5.2 - Collision 2.4.5.3 - Birthday
3.2	Cryptography Implementations	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.2 - Encryption

		1.4.2.1.1 - Full-disk 1.4.2.3 - Asymmetric 1.4.2.4 - Symmetric 1.4.2.5 - Key exchange • 1.4.3 - Tools 1.4.3.1 - Trusted Platform Module (TPM) 1.4.3.2 - Hardware security module (HSM) 1.4.3.3 - Key management system 1.4.3.4 - Secure enclave • 1.4.4 - Obfuscation 1.4.4.1 - Steganography • 1.4.5 - Hashing • 1.4.7 - Digital signatures
3.3	Hashing	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.2 - Encryption 1.4.2.6 - Algorithms • 1.4.5 - Hashing • 1.4.6 - Salting • 1.4.7 - Digital signatures
3.4	Encryption	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.2 - Encryption 1.4.2.1.1 - Full-disk 1.4.2.1.3 - File 1.4.2.1.5 - Database • 1.4.3 - Tools

		1.4.3.1 - Trusted Platform Module (TPM) 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.6 - Encryption
3.5	Public Key Infrastructure	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.1 - Public key infrastructure (PKI) 1.4.1.1 - Public key 1.4.1.2 - Private key 1.4.1.3 - Key escrow • 1.4.7 - Digital signatures • 1.4.11 - Certificates 1.4.11.1 - Certificate authorities 1.4.11.2 - Certificate revocation lists (CRLs) 1.4.11.3 - Online Certificate Status Protocol (OCSP) 1.4.11.4 - Self-signed 1.4.11.5 - Third-party 1.4.11.6 - Root of trust 1.4.11.7 - Certificate signing request (CSR) generation 1.4.11.8 - Wildcard 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.6 - Application security 4.1.6.1 - Input validation 4.1.6.2 - Secure cookies 4.1.6.3 - Static code analysis 4.1.6.4 - Code signing

4.0	Identity and Access Management	
4.1	Access Control Models	1.2 Summarize fundamental security concepts • 1.2.1 - Confidentiality, Integrity, and Availability (CIA) • 1.2.2 - Non-repudiation • 1.2.3 - Authentication, Authorization, and Accounting (AAA) 1.2.3.1 - Authenticating people 1.2.3.2 - Authenticating systems 1.2.3.3 - Authorization models • 1.2.4 - Gap analysis • 1.2.5 - Zero trust 1.2.5.1 - Control plane 1.2.5.1.1 - Adaptive identity 1.2.5.1.2 - Threat scope reduction 1.2.5.1.3 - Policy-driven access control 1.2.5.1.4 - Policy Administrator 1.2.5.1.5 - Policy Engine 1.2.5.2 - Data plane 1.2.5.2.1 - Implicit trust zones 1.2.5.2.2 - Subject/System 1.2.5.2.3 - Policy enforcement point 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.2 - Access control • 2.5.8 - Least privilege 4.6 Given a scenario, implement and maintain identity and access management • 4.6.1 - Provisioning/de-provisioning user accounts • 4.6.2 - Permission assignments and implications • 4.6.3 - Identity proofing • 4.6.8 - Access controls

		4.6.8.1 - Mandatory 4.6.8.2 - Discretionary 4.6.8.3 - Role-based 4.6.8.4 - Rule-based 4.6.8.5 - Attribute-based 4.6.8.6 - Time-of-day restrictions 4.6.8.7 - Least privilege • 4.6.9 - Multifactor authentication 4.6.9.1.2 - Hard/soft authentication tokens 4.6.9.2 - Factors 4.6.9.2.1 - Something you know 4.6.9.2.2 - Something you have 4.6.9.2.3 - Something you are 4.6.9.2.4 - Somewhere you are 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.2 - Information security policies
4.2	Authentication	1.2 Summarize fundamental security concepts • 1.2.3 - Authentication, Authorization, and Accounting (AAA) 1.2.3.1 - Authenticating people 1.2.3.2 - Authenticating systems 1.2.3.3 - Authorization models 4.6 Given a scenario, implement and maintain identity and access management • 4.6.4 - Federation • 4.6.5 - Single sign-on (SSO)

		4.6.5.1 - Lightweight Directory Access Protocol (LDAP) 4.6.5.2 - Open authorization (OAuth) 4.6.5.3 - Security Assertions Markup Language (SAML) • 4.6.6 - Interoperability • 4.6.7 - Attestation • 4.6.9 - Multifactor authentication 4.6.9.1 - Implementations 4.6.9.1.1 - Biometrics 4.6.9.1.2 - Hard/soft authentication tokens 4.6.9.1.3 - Security keys 4.6.9.2 - Factors 4.6.9.2.1 - Something you know 4.6.9.2.2 - Something you have 4.6.9.2.3 - Something you are 4.6.9.2.4 - Somewhere you are
4.3	Authorization	1.2 Summarize fundamental security concepts • 1.2.3 - Authentication, Authorization, and Accounting (AAA) 1.2.3.3 - Authorization models 1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.4 - Obfuscation 1.4.4.2 - Tokenization 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.2 - Access control 2.5.2.1 - Access control list (ACL) 2.5.2.2 - Permissions

		4.6 Given a scenario, implement and maintain identity and access management • 4.6.2 - Permission assignments and implications • 4.6.5 - Single sign-on (SSO) • 4.6.8 - Access controls 4.6.8.2 - Discretionary
4.4	Active Directory Overview	1.2 Summarize fundamental security concepts • 1.2.3 - Authentication, Authorization, and Accounting (AAA) 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security 4.5.4.1 - Group Policy
4.5	Hardening Authentication	1.2 Summarize fundamental security concepts • 1.2.3 - Authentication, Authorization, and Accounting (AAA) 1.2.3.1 - Authenticating people 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.7 - Indicators 2.4.7.1 - Account lockout 2.4.7.2 - Concurrent session usage 2.4.7.3 - Blocked content 2.4.7.4 - Impossible travel 2.4.7.6 - Resource inaccessibility 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security 4.5.4.1 - Group Policy 4.6 Given a scenario, implement and maintain identity and access management • 4.6.1 - Provisioning/de-provisioning user accounts

		• 4.6.2 - Permission assignments and implications • 4.6.9 - Multifactor authentication 4.6.9.1.3 - Security keys 4.6.9.2.2 - Something you have • 4.6.10 - Password concepts 4.6.10.1 - Password best practices 4.6.10.1.1 - Length 4.6.10.1.2 - Complexity 4.6.10.1.3 - Reuse 4.6.10.1.4 - Expiration 4.6.10.1.5 - Age 4.6.10.2 - Password managers 4.6.10.3 - Passwordless • 4.6.11 - Privileged access management tools 4.6.11.1 - Just-in-time permissions 4.6.11.2 - Password vaulting 4.6.11.3 - Temporal accounts
4.6	Linux Users	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.11 - Hardening techniques 2.5.11.6 - Default password changes 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.6 - Application security 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security

		4.5.4.2 - SELinux 4.6 Given a scenario, implement and maintain identity and access management • 4.6.1 - Provisioning/de-provisioning user accounts • 4.6.10 - Password concepts 4.6.10.1.3 - Reuse 4.6.10.1.4 - Expiration
4.7	Linux Groups	4.6 Given a scenario, implement and maintain identity and access management • 4.6.1 - Provisioning/de-provisioning user accounts
4.8	Remote Access	3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.1 - Virtual private network (VPN) 3.2.2.2 - Remote access 3.2.2.3 - Tunneling 3.2.2.3.2 - Internet protocol security (IPSec) 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.5 - Wireless security settings 4.1.5.2 - AAA/Remote Authentication Dial-In User Service (RADIUS)
4.9	Network Authentication	4.6 Given a scenario, implement and maintain identity and access management • 4.6.4 - Federation • 4.6.5 - Single sign-on (SSO) 4.6.5.1 - Lightweight Directory Access Protocol (LDAP) 4.6.5.2 - Open authorization (OAuth) 4.6.5.3 - Security Assertions Markup Language (SAML)
5.0	**Network Architecture**	

5.1	Enterprise Network Architecture	3.1 Compare and contrast security implications of different architecture models • 3.1.1 - Architecture and infrastructure concepts 3.1.1.1 - Cloud 3.1.1.1.1 - Responsibility matrix 3.1.1.1.2 - Hybrid considerations 3.1.1.1.3 - Third-party vendors 3.1.1.2 - Infrastructure as code (IaC) 3.1.1.3 - Serverless 3.1.1.4 - Microservices 3.1.1.5 - Network infrastructure 3.1.1.5.1 - Physical isolation 3.1.1.5.1.1 - Air-gapped 3.1.1.5.2 - Logical segmentation 3.1.1.5.3 - Software-defined networking (SDN) 3.1.1.6 - On-premises 3.1.1.7 - Centralized/decentralized • 3.1.2 - Considerations 3.1.2.1 - Availability 3.1.2.2 - Resilience 3.1.2.3 - Cost 3.1.2.4 - Responsiveness 3.1.2.5 - Scalability 3.1.2.6 - Ease of deployment 3.1.2.7 - Risk transference

		3.1.2.8 - Ease of recovery 3.1.2.9 - Patch availability 3.1.2.10 - Inability to patch 3.1.2.11 - Power 3.1.2.12 - Compute 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.1 - Device placement 3.2.1.2 - Security zones 3.2.1.3 - Attack surface 3.2.1.4 - Connectivity 3.2.1.5 - Failure modes 3.2.1.5.1 - Fail-open 3.2.1.5.2 - Fail-closed 3.2.1.6 - Device attribute 3.2.1.6.1 - Active vs. passive 3.2.1.6.2 - Inline vs. tap/monitor 3.2.1.7.4 - Load balancer • 3.2.3 - Selection of effective controls
5.2	Security Appliances	1.2 Summarize fundamental security concepts • 1.2.7 - Deception and disruption technology 1.2.7.1 - Honeypot 1.2.7.2 - Honeynet 1.2.7.3 - Honeyfile 1.2.7.4 - Honeytoken

		3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.2 - Security zones 3.2.1.5 - Failure modes 3.2.1.5.1 - Fail-open 3.2.1.5.2 - Fail-closed 3.2.1.7 - Network appliances 3.2.1.7.1 - Jump server 3.2.1.7.2 - Proxy server 3.2.1.7.3 - Intrusion prevention system (IPS)/intrusion detection system (IDS) 3.2.1.7.4 - Load balancer 3.2.1.7.5 - Sensors 3.2.1.9.2 - Unified threat management (UTM) 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.1 - Firewall 4.5.1.4 - Screened subnets • 4.5.2 - IDS/IPS 4.5.2.1 - Trends 4.5.2.2 - Signatures • 4.5.3 - Web filter 4.5.3.1 - Agent-based 4.5.3.2 - Centralized proxy 4.5.3.3 - Universal Resource Locator (URL) scanning

		4.5.3.4 - Content categorization 4.5.3.5 - Block rules 4.5.3.6 - Reputation • 4.5.6 - DNS filtering
5.3	Screened Subnets	3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.1 - Virtual private network (VPN) 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.1 - Firewall 4.5.1.4 - Screened subnets
5.4	Firewalls	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.11 - Hardening techniques 2.5.11.3 - Host-based firewall 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.9 - Firewall types 3.2.1.9.1 - Web application firewall (WAF) 3.2.1.9.2 - Unified threat management (UTM) 3.2.1.9.3 - Next-generation firewall (NGFW) 3.2.1.9.4 - Layer 4/Layer 7 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.1 - Firewall 4.5.1.1 - Rules

		4.5.1.2 - Access lists 4.5.1.3 - Ports/protocols 4.5.1.4 - Screened subnets
5.5	Virtual Private Networks	3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.1 - Virtual private network (VPN) 3.2.2.2 - Remote access 3.2.2.3 - Tunneling 3.2.2.3.1 - Transport Layer Security (TLS) 3.2.2.3.2 - Internet protocol security (IPSec) 3.2.2.4 - Software-defined wide area network (SD-WAN) 3.2.2.5 - Secure access service edge (SASE)
5.6	Network Access Control	4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.10. Network access control (NAC)
5.7	Network Device Vulnerabilities	2.3 Explain various types of vulnerabilities • 2.3.9 - Misconfiguration • 2.3.11 - Zero-day 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.4 - Application attacks 2.4.4.4 - Privilege escalation 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.1 - Segmentation • 2.5.11 - Hardening techniques 2.5.11.6 - Default password changes 4.6 Given a scenario, implement and maintain identity and access management

		• 4.6.1 - Provisioning/de-provisioning user accounts • 4.6.10 - Password concepts 4.6.10.1 - Password best practices 4.6.10.1.1 - Length 4.6.10.1.2 - Complexity 4.6.10.1.3 - Reuse 4.6.10.1.4 - Expiration 4.6.10.1.5 - Age 5.1 Summarize elements of effective security governance • 5.1.3 - Standards 5.1.3.1 - Password
5.8	Network Applications	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.7 - Digital signatures 4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods 4.3.1.1 - Vulnerability scan 4.3.1.2 - Application security 4.3.1.2.1 - Static analysis 4.3.1.2.2 - Dynamic analysis 4.3.1.2.3 - Package monitoring • 4.3.3 - Vulnerability response and remediation 4.3.3.1 - Patching
5.9	Switch Security and Attacks	3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations

		3.2.1.8 - Port security 3.2.1.8.1 - 802.1X 3.2.1.8.2 - Extensible Authentication Protocol (EAP) 3.3 Compare and contrast concepts and strategies to protect data • 3.3.4 - Methods to secure data 3.3.4.7 - Segmentation 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.2 - Hardening targets 4.1.2.3 - Switches • 4.1.5 - Wireless security settings 4.1.5.2 - AAA/Remote Authentication Dial-In User Service (RADIUS) 4.1.5.4 - Authentication protocols 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.5 - Implementation of secure protocols
5.10	Router Security	1.1 Compare and contrast various types of security controls • 1.1.1 - Categories 1.1.1.4 - Physical 2.2 Explain common threat vectors and attack surfaces • 2.2.10 - Default credentials 2.3 Explain various types of vulnerabilities • 2.3.4 - Hardware 2.3.4.1 - Firmware 2.5 Explain the purpose of mitigation techniques used to secure the enterprise

		• 2.5.2 - Access control 2.5.2.1 - Access control list (ACL) 2.5.2.2 - Permissions • 2.5.11 - Hardening techniques 2.5.11.5 - Disabling ports/protocols 2.5.11.6 - Default password changes 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.2 - Remote access 3.3 Compare and contrast concepts and strategies to protect data • 3.3.4 - Methods to secure data 3.3.4.7 - Segmentation 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.2 - Hardening targets 4.1.2.4 - Routers 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.5 - Implementation of secure protocols 4.5.5.1 - Protocol selection 5.1 Summarize elements of effective security governance • 5.1.3 - Standards 5.1.3.2 - Access control
6.0	**Resiliency and Site Security**	
6.1	Physical Threats	1.2 Summarize fundamental security concepts • 1.2.6 - Physical security

		1.2.6.1 - Bollards 1.2.6.2 - Access control vestibule 1.2.6.3 - Fencing 1.2.6.4 - Video surveillance 1.2.6.5 - Security guard 1.2.6.6 - Access badge 1.2.6.7 - Lighting 1.2.6.8 - Sensors 1.2.6.8.1 - Infrared 1.2.6.8.2 - Pressure 1.2.6.8.3 - Microwave 1.2.6.8.4 - Ultrasonic 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.2 - Physical attacks 2.4.2.1 - Brute force 2.4.2.2 - Radio frequency identification (RFID) cloning 2.4.2.3 - Environmental 3.4 Explain the importance of resilience and recovery in security architecture • 3.4.9 - Power 4.6 Given a scenario, implement and maintain identity and access management • 4.6.9 - Multifactor authentication 4.6.9.1.1 - Biometrics
6.2	Monitoring and Reconnaissance	2.2 Explain common threat vectors and attack surfaces • 2.2.8 - Unsecure networks • 2.2.9 - Open service ports

2.5 Explain the purpose of mitigation techniques used to secure the enterprise

- 2.5.7 - Monitoring

4.1 Given a scenario, apply common security techniques to computing resources

- 4.1.8. Monitoring

4.3 Explain various activities associated with vulnerability management

- 4.3.1 - Identification methods
    - 4.3.1.3.1 - Open-source intelligence (OSINT)

4.4 Explain security alerting and monitoring concepts and tools

- 4.4.1 - Monitoring computing resources
    - 4.4.1.1 - Systems
    - 4.4.1.3 - Infrastructure
- 4.4.3 - Tools
    - 4.4.3.6 - Vulnerability scanners

4.5 Given a scenario, modify enterprise capabilities to enhance security

- 4.5.1 - Firewall
    - 4.5.1.3 - Ports/protocols

4.9 Given a scenario, use data sources to support an investigation

- 4.9.2 - Data sources
    - 4.9.2.4 - Packet captures

5.4 Summarize elements of effective security compliance

- 5.4.3 - Compliance monitoring

5.5 Explain types and purposes of audits and assessments

- 5.5.2 - Internal

		5.5.2.1 - Compliance • 5.5.4 - Penetration testing 5.5.4.8 - Reconnaissance 5.5.4.8.1 - Passive 5.5.4.8.2 - Active
6.3	Intrusion Detection	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.11 - Hardening techniques 2.5.11.4 - Host-based intrusion prevention system (HIPS) 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.6.2 - Inline vs. tap/monitor 3.2.1.7.3 - Intrusion prevention system (IPS)/intrusion detection system (IDS) 3.2.1.7.5 - Sensors 4.3 Explain various activities associated with vulnerability management • 4.3.2 - Analysis 4.3.2.1 - Confirmation 4.3.2.1.1 - False positive 4.3.2.1.2 - False negative 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.2 - IDS/IPS 4.5.2.1 - Trends 4.5.2.2 - Signatures

		• 4.5.12. User behavior analytics 5.6 Given a scenario, implement security awareness practices • 5.6.2 - Anomalous behavior recognition 5.6.2.2 - Unexpected
6.4	Protocol Analyzers	2.4 Given a scenario, analyze indicators of malicious activity • 2.4.3 - Network attacks 2.4.3.4 - On-path 2.4.3.5 - Credential replay 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.11 - Hardening techniques 2.5.11.5 - Disabling ports/protocols 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.8. Monitoring 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.1 - Firewall 4.5.1.3 - Ports/protocols • 4.5.5 - Implementation of secure protocols 4.9 Given a scenario, use data sources to support an investigation • 4.9.2 - Data sources 4.9.2.4 - Packet captures
6.5	Analyzing Network Attacks	2.1 Compare and contrast common threat actors and motivations • 2.1.3 - Motivations 2.1.3.1 - Data exfiltration 2.4 Given a scenario, analyze indicators of malicious activity

		• 2.4.3 - Network attacks 2.4.3.1 - Distributed denial-of-service (DDoS) 2.4.3.1.1 - Amplified 2.4.3.1.2 - Reflected 2.4.3.2 - Domain Name System (DNS) attacks 2.4.3.3 - Wireless 2.4.3.4 - On-path 2.4.3.5 - Credential replay 2.4.3.6 - Malicious code • 2.4.4 - Application attacks 2.4.4.4 - Privilege escalation 4.7 Explain the importance of automation and orchestration related to secure operations • 4.7.1 - Use cases of automation and scripting 4.9 Given a scenario, use data sources to support an investigation • 4.9.2 - Data sources 4.9.2.4 - Packet captures 5.5 Explain types and purposes of audits and assessments • 5.5.4 - Penetration testing 5.5.4.8 - Reconnaissance
6.6	Analyzing Password Attacks	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.5 - Hashing • 1.4.6 - Salting 2.2 Explain common threat vectors and attack surfaces • 2.2.10 - Default credentials • 2.2.12 - Human vectors/social engineering

		2.4 Given a scenario, analyze indicators of malicious activity • 2.4.6 - Password attacks 2.4.6.1 - Spraying 2.4.6.2 - Brute force 5.6 Given a scenario, implement security awareness practices • 5.6.3 - User guidance and training 5.6.3.4 - Password management
7.0	**Vulnerability Management**	
7.1	Vulnerability Management	1.1 Compare and contrast various types of security controls • 1.1.2 - Control types 1.1.2.5 - Compensating 2.3 Explain various types of vulnerabilities • 2.3.2 - Operating system (OS)-based • 2.3.4 - Hardware 2.3.4.1 - Firmware 2.3.4.2 - End-of-life 2.3.4.3 - Legacy 4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods 4.3.1.1 - Vulnerability scan 4.3.1.3 - Threat feed 4.3.1.3.1 - Open-source intelligence (OSINT) 4.3.1.3.2 - Proprietary/third-party 4.3.1.3.3 - Information-sharing organization 4.3.1.3.4 - Dark web

		4.3.1.4 - Penetration testing 4.3.1.5 - Responsible disclosure program 4.3.1.5.1 - Bug bounty program 4.3.1.6 - System/process audit • 4.3.2 - Analysis 4.3.2.1 - Confirmation 4.3.2.2 - Prioritize 4.3.2.3 - Common Vulnerability Scoring System (CVSS) 4.3.2.4 - Common Vulnerability Enumeration (CVE) 4.3.2.5 - Vulnerability classification 4.3.2.6 - Exposure factor 4.3.2.7 - Environmental variables 4.3.2.8 - Industry/organizational impact 4.3.2.9 - Risk tolerance • 4.3.3 - Vulnerability response and remediation 4.3.3.1 - Patching 4.3.3.2 - Insurance 4.3.3.3 - Segmentation 4.3.3.4 - Compensating controls 4.3.3.5 - Exceptions and exemptions • 4.3.4 - Validation of remediation 4.3.4.1 - Rescanning 4.3.4.2 - Audit 4.3.4.3 - Verification 4.4 Explain security alerting and monitoring concepts and tools

		• 4.4.3 - Tools 4.4.3.6 - Vulnerability scanners 4.8 Explain appropriate incident response activities • 4.8.5 - Threat hunting 4.9 Given a scenario, use data sources to support an investigation • 4.9.2 - Data sources 4.9.2.1 - Vulnerability scans
7.2	Vulnerability Scanning	2.2 Explain common threat vectors and attack surfaces • 2.2.6 - Vulnerable software 4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods 4.3.1.1 - Vulnerability scan 4.3.1.2 - Application security 4.3.1.2.1 - Static analysis 4.3.1.2.2 - Dynamic analysis 4.3.1.2.3 - Package monitoring • 4.3.2 - Analysis 4.3.2.4 - Common Vulnerability Enumeration (CVE) 4.3.2.5 - Vulnerability classification 4.4 Explain security alerting and monitoring concepts and tools • 4.4.1 - Monitoring computing resources 4.4.1.1 - Systems 4.4.1.2 - Applications 4.4.1.3 - Infrastructure • 4.4.2 - Activities

		4.4.2.3 - Scanning 4.4.2.4 - Reporting • 4.4.3 - Tools 4.4.3.6 - Vulnerability scanners 4.9 Given a scenario, use data sources to support an investigation • 4.9.2 - Data sources 4.9.2.3 - Dashboards
7.3	Alerting and Monitoring	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.4 - Obfuscation 1.4.4.2 - Tokenization 1.4.4.3 - Data masking 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.7 - Monitoring 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.7.5 - Sensors 3.3 Compare and contrast concepts and strategies to protect data • 3.3.3 - General data considerations 3.3.3.1.1 - Data at rest 3.3.3.1.2 - Data in transit 3.3.3.1.3 - Data in use • 3.3.4 - Methods to secure data 3.3.4.2 - Encryption

		3.3.4.4 - Masking 3.3.4.5 - Tokenization 3.3.4.8 - Permission restrictions 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.8. Monitoring 4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods 4.3.1.1 - Vulnerability scan 4.4 Explain security alerting and monitoring concepts and tools • 4.4.1 - Monitoring computing resources 4.4.1.1 - Systems 4.4.1.3 - Infrastructure • 4.4.2 - Activities 4.4.2.1 - Log aggregation 4.4.2.2 - Alerting 4.4.2.4 - Reporting 4.4.2.5 - Archiving 4.4.2.6 - Alert response and remediation/validation 4.4.2.6.2 - Alert tuning • 4.4.3 - Tools 4.4.3.1 - Security Content Automation Protocol (SCAP) 4.4.3.2 - Benchmarks 4.4.3.3 - Agents/agentless

		4.4.3.3.1 - Security information and event management (SIEM) 4.4.3.3.2 - Antivirus 4.4.3.3.3 - Data loss prevention (DLP) 4.4.3.4 - Simple Network Management Protocol (SNMP) traps 4.4.3.5 - NetFlow 4.4.3.6 - Vulnerability scanners 4.9 Given a scenario, use data sources to support an investigation • 4.9.1 - Log data 4.9.1.1 - Firewall logs 4.9.1.3 - Endpoint logs 4.9.1.4 - OS-specific security logs 4.9.1.5 - IPS/IDS logs 4.9.1.6 - Network logs • 4.9.2 - Data sources 4.9.2.3 - Dashboards 5.1 Summarize elements of effective security governance • 5.1.4 - Procedures 5.1.4.3 - Playbooks 5.6 Given a scenario, implement security awareness practices • 5.6.4 - Reporting and monitoring
7.4	Penetration Testing	2.2 Explain common threat vectors and attack surfaces • 2.2.12 - Human vectors/social engineering 4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods

		4.3.1.3.1 - Open-source intelligence (OSINT) 4.3.1.4 - Penetration testing 4.3.1.5.1 - Bug bounty program 4.4 Explain security alerting and monitoring concepts and tools • 4.4.3 - Tools 4.4.3.6 - Vulnerability scanners 5.3 Explain the processes associated with third-party risk assessment and management • 5.3.3 - Agreement types 5.3.3.5 - Work order (WO)/statement of work (SOW) • 5.3.6 - Rules of engagement 5.5 Explain types and purposes of audits and assessments • 5.5.4 - Penetration testing 5.5.4.1 - Physical 5.5.4.2 - Offensive 5.5.4.3 - Defensive 5.5.4.4 - Integrated 5.5.4.5 - Known environment 5.5.4.6 - Partially known environment 5.5.4.7 - Unknown environment 5.5.4.8 - Reconnaissance 5.5.4.8.1 - Passive 5.5.4.8.2 - Active
8.0	**Network and Endpoint Security**	

8.1	Operating System Hardening	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.2 - Encryption 1.4.2.1.1 - Full-disk 2.2 Explain common threat vectors and attack surfaces • 2.2.5 - Removable device 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.2 - Access control • 2.5.3 - Application allow list • 2.5.5 - Patching • 2.5.6 - Encryption • 2.5.7 - Monitoring • 2.5.8 - Least privilege • 2.5.9 - Configuration enforcement • 2.5.10 - Decommissioning • 2.5.11 - Hardening techniques 2.5.11.2 - Installation of endpoint protection 2.5.11.3 - Host-based firewall 2.5.11.5 - Disabling ports/protocols 2.5.11.6 - Default password changes 2.5.11.7 - Removal of unnecessary software 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.1 - Secure baselines 4.1.1.1 - Establish 4.1.1.2 - Deploy 4.1.1.3 - Maintain • 4.1.2 - Hardening targets 4.1.2.2 - Workstations

		4.3 Explain various activities associated with vulnerability management • 4.3.3 - Vulnerability response and remediation 4.3.3.1 - Patching 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security 4.5.4.1 - Group Policy 4.6 Given a scenario, implement and maintain identity and access management • 4.6.1 - Provisioning/de-provisioning user accounts • 4.6.8 - Access controls 4.6.8.7 - Least privilege • 4.6.9 - Multifactor authentication 4.7 Explain the importance of automation and orchestration related to secure operations • 4.7.1 - Use cases of automation and scripting 4.7.1.4 - Security groups
8.2	File Server Security	1.2 Summarize fundamental security concepts • 1.2.6 - Physical security 1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.2 - Encryption 1.4.2.1.1 - Full-disk 1.4.2.1.3 - File 2.2 Explain common threat vectors and attack surfaces • 2.2.3 - File-based 2.5 Explain the purpose of mitigation techniques used to secure the enterprise

- 2.5.2 - Access control
    - 2.5.2.1 - Access control list (ACL)
    - 2.5.2.2 - Permissions
- 2.5.8 - Least privilege
- 2.5.11 - Hardening techniques
    - 2.5.11.7 - Removal of unnecessary software

3.2 Given a scenario, apply security principles to secure enterprise infrastructure

- 3.2.2 - Secure communication/access
    - 3.2.2.1 - Virtual private network (VPN)
        - 3.2.2.3.2 - Internet protocol security (IPSec)

4.1 Given a scenario, apply common security techniques to computing resources

- 4.1.2 - Hardening targets
    - 4.1.2.6 - Servers

4.2 Explain the security implications of proper hardware, software, and data asset management

- 4.2.1 - Acquisition/procurement process

4.3 Explain various activities associated with vulnerability management

- 4.3.1 - Identification methods
    - 4.3.1.6 - System/process audit

4.5 Given a scenario, modify enterprise capabilities to enhance security

- 4.5.5 - Implementation of secure protocols
- 4.5.8. File integrity monitoring

4.6 Given a scenario, implement and maintain identity and access management

- 4.6.1 - Provisioning/de-provisioning user accounts
- 4.6.8 - Access controls

		4.6.8.7 - Least privilege
8.3	Linux Host Security	2.2 Explain common threat vectors and attack surfaces • 2.2.5 - Removable device 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.5 - Patching • 2.5.11 - Hardening techniques 2.5.11.3 - Host-based firewall 2.5.11.4 - Host-based intrusion prevention system (HIPS) 2.5.11.5 - Disabling ports/protocols 2.5.11.7 - Removal of unnecessary software
8.4	Wireless Overview	4.1 Given a scenario, apply common security techniques to computing resources • 4.1.3. Wireless devices 4.1.3.1 - Installation considerations 4.1.3.1.1 - Site surveys 4.1.3.1.2 - Heat maps • 4.1.5 - Wireless security settings 4.1.5.1 - Wi-Fi Protected Access 3 (WPA3) 4.1.5.2 - AAA/Remote Authentication Dial-In User Service (RADIUS)
8.5	Wireless Attacks	2.2 Explain common threat vectors and attack surfaces • 2.2.8 - Unsecure networks 2.2.8.1 - Wireless 2.2.8.3 - Bluetooth 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.2 - Physical attacks

		2.4.2.2 - Radio frequency identification (RFID) cloning • 2.4.3 - Network attacks 2.4.3.3 - Wireless 2.4.3.5 - Credential replay 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.3. Wireless devices 4.1.3.1 - Installation considerations • 4.1.5 - Wireless security settings 4.1.5.1 - Wi-Fi Protected Access 3 (WPA3) 4.1.5.4 - Authentication protocols
8.6	Wireless Defenses	2.2 Explain common threat vectors and attack surfaces • 2.2.8 - Unsecure networks 2.2.8.1 - Wireless • 2.2.10 - Default credentials 2.3 Explain various types of vulnerabilities • 2.3.4 - Hardware 2.3.4.1 - Firmware 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.11 - Hardening techniques 2.5.11.6 - Default password changes 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.7.3 - Intrusion prevention system (IPS)/intrusion detection system (IDS)

		3.2.1.8 - Port security 3.2.1.8.1 - 802.1X 3.2.1.8.2 - Extensible Authentication Protocol (EAP) • 3.2.2 - Secure communication/access 3.2.2.3.1 - Transport Layer Security (TLS) 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.3. Wireless devices 4.1.3.1 - Installation considerations • 4.1.5 - Wireless security settings 4.1.5.1 - Wi-Fi Protected Access 3 (WPA3) 4.1.5.2 - AAA/Remote Authentication Dial-In User Service (RADIUS) 4.1.5.3 - Cryptographic protocols 4.1.5.4 - Authentication protocols
8.7	Data Transmission Security	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.2 - Encryption 1.4.2.2 - Transport/communication 1.4.2.3 - Asymmetric 1.4.2.5 - Key exchange • 1.4.11 - Certificates 1.4.11.1 - Certificate authorities 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.3 - Network attacks 2.4.3.2 - Domain Name System (DNS) attacks

		3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.1 - Virtual private network (VPN) 3.2.2.2 - Remote access 3.2.2.3 - Tunneling 3.2.2.3.1 - Transport Layer Security (TLS) 3.2.2.3.2 - Internet protocol security (IPSec)
8.8	Web Application Security	2.3 Explain various types of vulnerabilities • 2.3.1 - Application 2.3.1.1 - Memory injection 2.3.1.2 - Buffer overflow 2.3.1.3 - Race conditions 2.3.1.3.1 - Time-of-check (TOC) 2.3.1.3.2 - Time-of-use (TOU) 2.3.1.4 - Malicious update • 2.3.3 - Web-based 2.3.3.1 - Structured Query Language injection (SQLi) 2.3.3.2 - Cross-site scripting (XSS) • 2.3.4 - Hardware • 2.3.11 - Zero-day 2.4 Given a scenario, analyze indicators of malicious activity • 2.4.3 - Network attacks 2.4.3.5 - Credential replay • 2.4.4 - Application attacks 2.4.4.1 - Injection

		2.4.4.2 - Buffer overflow 2.4.4.3 - Replay 2.4.4.4 - Privilege escalation 2.4.4.5 - Forgery 2.4.4.6 - Directory traversal 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.5 - Patching 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.7.3 - Intrusion prevention system (IPS)/intrusion detection system (IDS) 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.6 - Application security 4.1.6.2 - Secure cookies 4.6 Given a scenario, implement and maintain identity and access management • 4.6.5 - Single sign-on (SSO) 4.6.5.1 - Lightweight Directory Access Protocol (LDAP)
8.9	Application Development and Security	1.3 Explain the importance of change management processes and the impact to security • 1.3.4 -Version control 2.2 Explain common threat vectors and attack surfaces • 2.2.7 - Unsupported systems and applications 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.2 - Access control

2.5.2.2 - Permissions

- 2.5.3 - Application allow list
- 2.5.5 - Patching
- 2.5.7 - Monitoring
- 2.5.11 - Hardening techniques

2.5.11.2 - Installation of endpoint protection

2.5.11.3 - Host-based firewall

2.5.11.7 - Removal of unnecessary software

3.3 Compare and contrast concepts and strategies to protect data

- 3.3.4 - Methods to secure data

3.3.4.6 - Obfuscation

4.1 Given a scenario, apply common security techniques to computing resources

- 4.1.1 - Secure baselines

4.1.1.1 - Establish

- 4.1.6 - Application security

4.1.6.1 - Input validation

4.1.6.2 - Secure cookies

4.1.6.3 - Static code analysis

4.1.6.4 - Code signing

- 4.1.7. Sandboxing

4.3 Explain various activities associated with vulnerability management

- 4.3.1 - Identification methods

4.3.1.2 - Application security

4.3.1.2.1 - Static analysis

4.3.1.2.2 - Dynamic analysis

		4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security 4.5.4.2 - SELinux 4.6 Given a scenario, implement and maintain identity and access management • 4.6.1 - Provisioning/de-provisioning user accounts • 4.6.8 - Access controls 4.6.8.7 - Least privilege 4.7 Explain the importance of automation and orchestration related to secure operations • 4.7.1 - Use cases of automation and scripting 4.7.1.1 - User provisioning 4.7.1.2 - Resource provisioning 4.7.1.4 - Security groups 4.7.1.5 - Ticket creation 4.7.1.7 - Enabling/disabling services and access 4.7.1.9 - Integrations and Application programming interfaces (APIs) • 4.7.2 - Benefits 4.7.2.1 - Efficiency/time saving 4.7.2.2 - Enforcing baselines 4.7.2.5 - Staff retention • 4.7.3 - Other considerations 4.7.3.1 - Complexity 4.7.3.2 - Cost 4.7.3.3 - Single point of failure

		4.7.3.4 - Technical debt 4.7.3.5 - Ongoing supportability 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.6 - Software development lifecycle (SDLC) 5.2 Explain elements of the risk management process • 5.2.3 - Risk analysis 5.2.3.8 - Exposure factor
9.0	**Incident Response**	
9.1	Incident Response and Mitigation	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.1 - Segmentation • 2.5.4 - Isolation 4.8 Explain appropriate incident response activities • 4.8.1 - Process 4.8.1.1 - Preparation 4.8.1.2 - Detection 4.8.1.3 - Analysis 4.8.1.4 - Containment 4.8.1.5 - Eradication 4.8.1.6 - Recovery 4.8.1.7 - Lessons learned 5.1 Summarize elements of effective security governance • 5.1.4 - Procedures 5.1.4.3 - Playbooks
9.2	Log Management	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.7 - Monitoring

3.2 Given a scenario, apply security principles to secure enterprise infrastructure

- 3.2.1 - Infrastructure considerations

    3.2.1.7.3 - Intrusion prevention system (IPS)/intrusion detection system (IDS)

    3.2.1.7.5 - Sensors

4.1 Given a scenario, apply common security techniques to computing resources

- 4.1.1 - Secure baselines

    4.1.1.1 - Establish

4.3 Explain various activities associated with vulnerability management

- 4.3.1 - Identification methods

    4.3.1.1 - Vulnerability scan

- 4.3.2 - Analysis

    4.3.2.4 - Common Vulnerability Enumeration (CVE)

4.4 Explain security alerting and monitoring concepts and tools

- 4.4.2 - Activities

    4.4.2.1 - Log aggregation

    4.4.2.2 - Alerting

- 4.4.3 - Tools

    4.4.3.3.1 - Security information and event management (SIEM)

    4.4.3.5 - NetFlow

    4.4.3.6 - Vulnerability scanners

4.9 Given a scenario, use data sources to support an investigation

- 4.9.1 - Log data

		4.9.1.1 - Firewall logs 4.9.1.2 - Application logs 4.9.1.3 - Endpoint logs 4.9.1.4 - OS-specific security logs 4.9.1.5 - IPS/IDS logs 4.9.1.6 - Network logs 4.9.1.7 - Metadata • 4.9.2 - Data sources 4.9.2.3 - Dashboards 4.9.2.4 - Packet captures
9.3	Digital Forensics	4.8 Explain appropriate incident response activities • 4.8.6 - Digital forensics 4.8.6.1 - Legal hold 4.8.6.2 - Chain of custody 4.8.6.3 - Acquisition 4.8.6.4 - Reporting 4.8.6.5 - Preservation 4.8.6.6 - E-discovery 4.9 Given a scenario, use data sources to support an investigation • 4.9.2 - Data sources 4.9.2.1 - Vulnerability scans 4.9.2.3 - Dashboards 4.9.2.4 - Packet captures
9.4	Redundancy	3.1 Compare and contrast security implications of different architecture models • 3.1.1 - Architecture and infrastructure concepts

		3.1.1.1.3 - Third-party vendors 3.1.1.5 - Network infrastructure 3.1.1.14 - High availability • 3.1.2 - Considerations 3.1.2.11 - Power 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.6.1 - Active vs. passive 3.2.1.7.4 - Load balancer 3.4 Explain the importance of resilience and recovery in security architecture • 3.4.1 - High availability 3.4.1.1 - Load balancing vs. clustering • 3.4.2 - Site considerations 3.4.2.4 - Geographic dispersion • 3.4.9 - Power 3.4.9.1 - Generators 3.4.9.2 - Uninterruptible power supply (UPS) 5.3 Explain the processes associated with third-party risk assessment and management • 5.3.2 - Vendor selection
9.5	Backup and Restore	3.4 Explain the importance of resilience and recovery in security architecture • 3.4.8 - Backups 3.4.8.1 - Onsite/offsite 3.4.8.2 - Frequency

		3.4.8.3 - Encryption 3.4.8.4 - Snapshots 3.4.8.5 - Recovery 3.4.8.6 - Replication 3.4.8.7 - Journaling
10.0	**Protocol, App, and Cloud Security**	
10.1	Host Virtualization	2.3 Explain various types of vulnerabilities • 2.3.5 - Virtualization 2.3.5.1 - Virtual machine (VM) escape 2.3.5.2 - Resource reuse 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.1 - Segmentation 3.1 Compare and contrast security implications of different architecture models • 3.1.1 - Architecture and infrastructure concepts 3.1.1.5 - Network infrastructure 3.1.1.5.2 - Logical segmentation 3.1.1.8 - Containerization 3.1.1.9 - Virtualization 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.1 - Infrastructure considerations 3.2.1.7.4 - Load balancer 3.4 Explain the importance of resilience and recovery in security architecture • 3.4.1 - High availability

		3.4.1.1 - Load balancing vs. clustering 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.7. Sandboxing 4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods 4.3.1.1 - Vulnerability scan
10.2	Virtual Networking	3.1 Compare and contrast security implications of different architecture models • 3.1.1 - Architecture and infrastructure concepts 3.1.1.5 - Network infrastructure 3.1.1.9 - Virtualization 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.1 - Virtual private network (VPN) 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.7. Sandboxing
10.3	Software-Defined Networking	3.1 Compare and contrast security implications of different architecture models • 3.1.1 - Architecture and infrastructure concepts 3.1.1.5.3 - Software-defined networking (SDN)
10.4	Cloud Services	1.2 Summarize fundamental security concepts • 1.2.3 - Authentication, Authorization, and Accounting (AAA) 1.2.3.1 - Authenticating people 2.2 Explain common threat vectors and attack surfaces • 2.2.11 - Supply chain

2.2.11.1 - Managed service providers (MSPs)

2.2.11.2 - Vendors

2.2.11.3 - Suppliers

2.3 Explain various types of vulnerabilities

- 2.3.6 - Cloud-specific
- 2.3.7 - Supply chain

2.3.7.1 - Service provider

3.1 Compare and contrast security implications of different architecture models

- 3.1.1 - Architecture and infrastructure concepts

3.1.1.1 - Cloud

3.1.1.1.2 - Hybrid considerations

3.1.1.1.3 - Third-party vendors

3.1.1.3 - Serverless

3.1.1.9 - Virtualization

- 3.1.2 - Considerations

3.4 Explain the importance of resilience and recovery in security architecture

- 3.4.4 - Multi-cloud systems

4.1 Given a scenario, apply common security techniques to computing resources

- 4.1.2 - Hardening targets

4.1.2.5 - Cloud infrastructure

5.3 Explain the processes associated with third-party risk assessment and management

- 5.3.1 - Vendor assessment

5.3.1.1 - Penetration testing

		5.3.1.2 - Right-to-audit clause 5.3.1.4 - Independent assessments • 5.3.4 - Vendor monitoring
10.5	Mobile Devices	2.3 Explain various types of vulnerabilities • 2.3.10 - Mobile device 2.3.10.1 - Side loading 2.3.10.2 - Jailbreaking 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.3 - Application allow list • 2.5.11 - Hardening techniques 2.5.11.1 - Encryption 3.3 Compare and contrast concepts and strategies to protect data • 3.3.4 - Methods to secure data 3.3.4.2 - Encryption 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.1 - Secure baselines 4.1.1.1 - Establish 4.1.1.2 - Deploy • 4.1.2 - Hardening targets 4.1.2.1 - Mobile devices • 4.1.3. Wireless devices 4.1.3.1 - Installation considerations • 4.1.4 - Mobile solutions 4.1.4.1 - Mobile device management (MDM) 4.1.4.2 - Deployment models

		4.1.4.2.1 - Bring your own device (BYOD) 4.1.4.2.2 - Corporate-owned, personally enabled (COPE) 4.1.4.2.3 - Choose your own device (CYOD) 4.1.4.3 - Connections methods 4.1.4.3.1 - Cellular 4.1.4.3.2 - Wi-Fi 4.1.4.3.3 - Bluetooth • 4.1.6 - Application security 4.1.6.4 - Code signing • 4.1.7. Sandboxing 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security 4.5.4.1 - Group Policy 4.6 Given a scenario, implement and maintain identity and access management • 4.6.10 - Password concepts 4.6.10.2 - Password managers
10.6	Mobile Device Management	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.3 - Application allow list 3.1 Compare and contrast security implications of different architecture models • 3.1.2 - Considerations 3.1.2.6 - Ease of deployment 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.4 - Mobile solutions

		4.1.4.1 - Mobile device management (MDM) 4.1.4.2 - Deployment models 4.1.4.2.1 - Bring your own device (BYOD) 4.2 Explain the security implications of proper hardware, software, and data asset management • 4.2.3 - Disposal/decommissioning 4.2.3.1 - Sanitization 4.2.3.2 - Destruction 4.2.3.4 - Data retention
10.7	BYOD Security	2.2 Explain common threat vectors and attack surfaces • 2.2.8 - Unsecure networks 2.2.8.1 - Wireless 2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.4 - Isolation 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.2 - Hardening targets 4.1.2.1 - Mobile devices • 4.1.3. Wireless devices 4.1.3.1 - Installation considerations • 4.1.4 - Mobile solutions 4.1.4.1 - Mobile device management (MDM) 4.1.4.2 - Deployment models 4.1.4.2.1 - Bring your own device (BYOD) 4.1.4.2.2 - Corporate-owned, personally enabled (COPE)

		4.1.4.2.3 - Choose your own device (CYOD) 4.1.4.3.2 - Wi-Fi 4.1.4.3.3 - Bluetooth 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.10. Network access control (NAC) 4.6 Given a scenario, implement and maintain identity and access management • 4.6.8 - Access controls 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.1 - Acceptable use policy (AUP)
10.8	Embedded and Specialized Systems	1.4 Explain the importance of using appropriate cryptographic solutions • 1.4.3 - Tools 1.4.3.1 - Trusted Platform Module (TPM) 2.2 Explain common threat vectors and attack surfaces • 2.2.8 - Unsecure networks 2.2.8.1 - Wireless 2.2.8.2 - Wired 2.2.8.3 - Bluetooth 3.1 Compare and contrast security implications of different architecture models • 3.1.1 - Architecture and infrastructure concepts 3.1.1.10 - IoT 3.1.1.11 - Industrial control systems (ICS)/supervisory control and data acquisition (SCADA) 3.1.1.12 - Real-time operating system (RTOS)

		3.1.1.13 - Embedded systems 3.3 Compare and contrast concepts and strategies to protect data • 3.3.4 - Methods to secure data 3.3.4.2 - Encryption 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.2 - Hardening targets 4.1.2.1 - Mobile devices 4.1.2.7 - ICS/SCADA 4.1.2.8 - Embedded systems 4.1.2.10 - IoT devices • 4.1.4 - Mobile solutions 4.1.4.3.1 - Cellular 4.1.4.3.2 - Wi-Fi 4.1.4.3.3 - Bluetooth
10.9	Email	2.4 Given a scenario, analyze indicators of malicious activity • 2.4.1 - Malware attacks 2.4.1.6 - Virus • 2.4.4 - Application attacks 3.2 Given a scenario, apply security principles to secure enterprise infrastructure • 3.2.2 - Secure communication/access 3.2.2.3.1 - Transport Layer Security (TLS) 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.6 - Application security 4.1.6.2 - Secure cookies

		4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.7 - Email security 4.5.7.1 - Domain-based Message Authentication Reporting and Conformance (DMARC) 4.5.7.2 - DomainKeys Identified Mail (DKIM) 4.5.7.3 - Sender Policy Framework (SPF) 4.5.7.4 - Gateway • 4.5.9. DLP 5.6 Given a scenario, implement security awareness practices • 5.6.1 - Phishing
11.0	**Security Governance Concepts**	
11.1	Policies, Standards, and Procedures	4.6 Given a scenario, implement and maintain identity and access management • 4.6.9 - Multifactor authentication • 4.6.10 - Password concepts 4.6.10.1.1 - Length 4.6.10.1.2 - Complexity 4.6.10.1.3 - Reuse 4.6.10.1.4 - Expiration 5.1 Summarize elements of effective security governance • 5.1.1 - Guidelines • 5.1.2 - Policies 5.1.2.1 - Acceptable use policy (AUP) 5.1.2.2 - Information security policies 5.1.2.3 - Business continuity 5.1.2.4 - Disaster recovery

		5.1.2.5 - Incident response 5.1.2.6 - Software development lifecycle (SDLC) 5.1.2.7 - Change management • 5.1.3 - Standards 5.1.3.1 - Password 5.1.3.2 - Access control 5.1.3.3 - Physical security 5.1.3.4 - Encryption • 5.1.4 - Procedures 5.1.4.3 - Playbooks • 5.1.5 - External considerations 5.1.5.1 - Regulatory 5.1.5.2 - Legal 5.1.5.3 - Industry 5.1.5.4 - Local/regional 5.1.5.5 - National 5.1.5.6 - Global • 5.1.6 - Monitoring and revision • 5.1.7 - Types of governance structures 5.1.7.1 - Boards 5.1.7.2 - Committees 5.1.7.3 - Government entities 5.1.7.4 - Centralized/decentralized
11.2	Change Management	1.3 Explain the importance of change management processes and the impact to security • 1.3.2 - Technical implications

		1.3.2.1 - Allow lists/deny lists 1.3.2.2 - Restricted activities 1.3.2.3 - Downtime 1.3.2.4 - Service restart 1.3.2.5 - Application restart 1.3.2.6 - Legacy applications 1.3.2.7 - Dependencies • 1.3.3 - Documentation 1.3.3.1 - Updating diagrams 1.3.3.2 - Updating policies/procedures • 1.3.4 -Version control 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.7 - Change management • 5.1.4 - Procedures 5.1.4.1 - Change management • 5.1.7 - Types of governance structures 5.1.7.1 - Boards 5.2 Explain elements of the risk management process • 5.2.3 - Risk analysis
11.3	Automation and Orchestration	4.7 Explain the importance of automation and orchestration related to secure operations • 4.7.1 - Use cases of automation and scripting 4.7.1.5 - Ticket creation 4.7.1.7 - Enabling/disabling services and access

		4.7.1.8 - Continuous integration and testing • 4.7.2 - Benefits 4.7.2.1 - Efficiency/time saving 4.7.2.2 - Enforcing baselines 4.7.2.3 - Standard infrastructure configurations 4.7.2.5 - Staff retention 4.7.2.6 - Reaction time 4.7.2.7 - Workforce multiplier • 4.7.3 - Other considerations 4.7.3.1 - Complexity 4.7.3.2 - Cost 4.7.3.3 - Single point of failure 4.7.3.4 - Technical debt 4.7.3.5 - Ongoing supportability
12.0	**Risk Management Processes**	
12.1	Risk Management Processes and Concepts	3.4 Explain the importance of resilience and recovery in security architecture • 3.4.1 - High availability • 3.4.2 - Site considerations 3.4.2.1 - Hot 3.4.2.2 - Cold 3.4.2.3 - Warm 3.4.2.4 - Geographic dispersion • 3.4.3 - Platform diversity • 3.4.4 - Multi-cloud systems • 3.4.5 - Continuity of operations • 3.4.6 - Capacity planning

		3.4.6.1 - People 3.4.6.2 - Technology 3.4.6.3 - Infrastructure • 3.4.7 - Testing 3.4.7.1 - Tabletop exercises 3.4.7.2 - Fail over 3.4.7.3 - Simulation 3.4.7.4 - Parallel processing • 3.4.8 - Backups 4.2 Explain the security implications of proper hardware, software, and data asset management • 4.2.2 - Assignment/accounting 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.3 - Business continuity 5.1.2.5 - Incident response 5.2 Explain elements of the risk management process • 5.2.1 - Risk identification • 5.2.2 - Risk assessment 5.2.2.1 - Ad hoc 5.2.2.2 - Recurring 5.2.2.3 - One-time 5.2.2.4 - Continuous • 5.2.3 - Risk analysis 5.2.3.1 - Qualitative

		5.2.3.2 - Quantitative 5.2.3.3 - Single loss expectancy (SLE) 5.2.3.4 - Annualized loss expectancy (ALE) 5.2.3.5 - Annualized rate of occurrence (ARO) 5.2.3.6 - Probability 5.2.3.7 - Likelihood 5.2.3.8 - Exposure factor 5.2.3.9 - Impact • 5.2.4 - Risk register 5.2.4.1 - Key risk indicators 5.2.4.2 - Risk owners 5.2.4.3 - Risk threshold • 5.2.5 - Risk tolerance • 5.2.6 - Risk appetite 5.2.6.1 - Expansionary 5.2.6.2 - Conservative 5.2.6.3 - Neutral • 5.2.7 - Risk management strategies 5.2.7.1 - Transfer 5.2.7.2 - Accept 5.2.7.2.1 - Exemption 5.2.7.2.2 - Exception 5.2.7.3 - Avoid 5.2.7.4 - Mitigate • 5.2.8 - Risk reporting • 5.2.9 - Business impact analysis

		5.2.9.1 - Recovery time objective (RTO) 5.2.9.2 - Recovery point objective (RPO) 5.2.9.3 - Mean time to repair (MTTR) 5.2.9.4 - Mean time between failures (MTBF)
12.2	Vendor Management	5.1 Summarize elements of effective security governance • 5.1.4 - Procedures 5.1.4.2 - Onboarding/offboarding 5.3 Explain the processes associated with third-party risk assessment and management • 5.3.1 - Vendor assessment 5.3.1.1 - Penetration testing 5.3.1.2 - Right-to-audit clause 5.3.1.3 - Evidence of internal audits 5.3.1.4 - Independent assessments 5.3.1.5 - Supply chain analysis • 5.3.2 - Vendor selection 5.3.2.1 - Due diligence 5.3.2.2 - Conflict of interest • 5.3.3 - Agreement types 5.3.3.1 - Service-level agreement (SLA) 5.3.3.2 - Memorandum of agreement (MOA) 5.3.3.3 - Memorandum of understanding (MOU) 5.3.3.4 - Master service agreement (MSA) 5.3.3.5 - Work order (WO)/statement of work (SOW) 5.3.3.6 - Non-disclosure agreement (NDA)

		5.3.3.7 - Business partners agreement (BPA) • 5.3.4 - Vendor monitoring • 5.3.5 - Questionnaires • 5.3.6 - Rules of engagement
12.3	Audits and Assessments	4.3 Explain various activities associated with vulnerability management • 4.3.1 - Identification methods 4.3.1.6 - System/process audit • 4.3.4 - Validation of remediation 4.3.4.2 - Audit 4.4 Explain security alerting and monitoring concepts and tools • 4.4.2 - Activities 4.4.2.1 - Log aggregation 4.5 Given a scenario, modify enterprise capabilities to enhance security • 4.5.4 - Operating system security 4.5.4.1 - Group Policy 4.9 Given a scenario, use data sources to support an investigation • 4.9.1 - Log data 4.9.1.4 - OS-specific security logs 4.9.1.6 - Network logs 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.2 - Information security policies 5.5 Explain types and purposes of audits and assessments • 5.5.1 - Attestation • 5.5.2 - Internal 5.5.2.1 - Compliance

		5.5.2.2 - Audit committee 5.5.2.3 - Self-assessments • 5.5.3 - External 5.5.3.1 - Regulatory 5.5.3.2 - Examinations 5.5.3.3 - Assessment 5.5.3.4 - Independent third-party audit
13.0	**Data Protection and Compliance**	
13.1	Data Classification and Compliance	3.3 Compare and contrast concepts and strategies to protect data • 3.3.1 - Data types 3.3.1.1 - Regulated 3.3.1.2 - Trade secret 3.3.1.3 - Intellectual property 3.3.1.4 - Legal information 3.3.1.5 - Financial information 3.3.1.6 - Human and non-human readable • 3.3.2 - Data classifications 3.3.2.1 - Sensitive 3.3.2.2 - Confidential 3.3.2.3 - Public 3.3.2.4 - Restricted 3.3.2.5 - Private 3.3.2.6 - Critical • 3.3.3 - General data considerations 3.3.3.2 - Data sovereignty

4.2 Explain the security implications of proper hardware, software, and data asset management

- 4.2.3 - Disposal/decommissioning
    - 4.2.3.1 - Sanitization
    - 4.2.3.2 - Destruction
    - 4.2.3.3 - Certification
    - 4.2.3.4 - Data retention

5.1 Summarize elements of effective security governance

- 5.1.8 - Roles and responsibilities for systems and data
    - 5.1.8.2 - Controllers
    - 5.1.8.3 - Processors

5.4 Summarize elements of effective security compliance

- 5.4.2 - Consequences of non-compliance
    - 5.4.2.1 - Fines
    - 5.4.2.2 - Sanctions
    - 5.4.2.3 - Reputational damage
    - 5.4.2.4 - Loss of license
    - 5.4.2.5 - Contractual impacts
- 5.4.4 - Privacy
    - 5.4.4.1 - Legal implications
        - 5.4.4.1.1 - Local/regional
        - 5.4.4.1.2 - National
        - 5.4.4.1.3 - Global
    - 5.4.4.2 - Data subject
    - 5.4.4.3 - Controller vs. processor

		5.4.4.4 - Ownership 5.4.4.5 - Data inventory and retention 5.4.4.6 - Right to be forgotten
13.2	Personnel Policies	2.5 Explain the purpose of mitigation techniques used to secure the enterprise • 2.5.8 - Least privilege 3.3 Compare and contrast concepts and strategies to protect data • 3.3.1 - Data types 3.3.1.1 - Regulated 3.3.1.3 - Intellectual property 3.3.1.4 - Legal information 3.3.1.5 - Financial information • 3.3.2 - Data classifications • 3.3.3 - General data considerations 3.3.3.1 - Data states 3.3.3.1.1 - Data at rest 3.3.3.1.2 - Data in transit 3.3.3.1.3 - Data in use 3.3.3.2 - Data sovereignty 3.3.3.3 - Geolocation • 3.3.4 - Methods to secure data 3.3.4.2 - Encryption 4.1 Given a scenario, apply common security techniques to computing resources • 4.1.4 - Mobile solutions 4.1.4.2.1 - Bring your own device (BYOD)

		4.2 Explain the security implications of proper hardware, software, and data asset management • 4.2.3 - Disposal/decommissioning 4.2.3.4 - Data retention 4.4 Explain security alerting and monitoring concepts and tools • 4.4.3 - Tools 4.4.3.3.3 - Data loss prevention (DLP) 4.6 Given a scenario, implement and maintain identity and access management • 4.6.8 - Access controls 4.6.8.7 - Least privilege 4.8 Explain appropriate incident response activities • 4.8.2 - Training • 4.8.6 - Digital forensics 4.8.6.1 - Legal hold 5.1 Summarize elements of effective security governance • 5.1.2 - Policies 5.1.2.1 - Acceptable use policy (AUP) • 5.1.4 - Procedures 5.1.4.2 - Onboarding/offboarding 5.3 Explain the processes associated with third-party risk assessment and management • 5.3.2 - Vendor selection 5.3.2.1 - Due diligence 5.4 Summarize elements of effective security compliance • 5.4.4 - Privacy 5.4.4.1 - Legal implications

		5.4.4.5 - Data inventory and retention 5.6 Given a scenario, implement security awareness practices • 5.6.1 - Phishing 5.6.1.2 - Recognizing a phishing attempt • 5.6.2 - Anomalous behavior recognition 5.6.2.1 - Risky 5.6.2.2 - Unexpected 5.6.2.3 - Unintentional • 5.6.3 - User guidance and training 5.6.3.1 - Policy/handbooks 5.6.3.2 - Situational awareness 5.6.3.5 - Removable media and cables 5.6.3.6 - Social engineering • 5.6.4 - Reporting and monitoring 5.6.4.1 - Initial 5.6.4.2 - Recurring • 5.6.5 - Development
A.0	**CompTIA Security+ SY0-701 - Practice Exams**	
A.1	Prepare for CompTIA Security+ SY0-701 Certification	
A.2	CompTIA Security+ Domain Review (20 Questions)	
A.3	CompTIA Security+ Domain Review (All Questions)	
B.0	**TestOut Security Pro - Practice Exams**	
B.1	Prepare for TestOut Security Pro Certification	

B.2	TestOut Security Pro Domain Review	

A.1.3 How to take the Security+ SY0-701 Exam

The Security+ SY0-701 exam is scheduled through Pearson VUE.

Certification	Provider	Website	Customer Service
CompTIA	Pearson VUE	pearsonvue.com	Online Chat and phone numbers by region are available on pearsonvue.com/comptia/contact

The Security+ SY0-701 exam information web page on CompTIA's website provides the latest details on how to schedule the exam.

TestOut is pleased to offer students a 10% off coupon code for exam vouchers purchased from CompTIA's online marketplace.

To purchase a voucher, go to the CompTIA website and complete the following:

Select your desired ***certification*** .

Add the corresponding ***CompTIA voucher(s)*** to your shopping cart.

Enter the coupon code **TESTOUT10** at checkout to get 10% off your purchase.

Go to the Pearson VUE website.

Enter your ***voucher information*** to register for the certification.

A.1.4 Security+ SY0-701 FAQs

Where do I take an exam?

When you schedule the exam with Pearson VUE, you will be presented with options to take the exam at a local testing center or your home or office. If you choose a local testing center, you will be given the option to pick from the testing centers closest to your location.

What do I bring to the exam?

You will need two forms of identification, one with a picture. For example, you could use a driver's license and a credit card. You will typically receive an erasable marker and whiteboard or laminated paper to use during the exam. Notes or other reference materials are not allowed inside the testing center. It is recommended that you do not bring personal items to the testing center; however, a testing center typically provides a locker to store personal items.

What specific information do I need to know to pass the exam?

People often ask, "What is on the exam?" This course is intended to help you gain the knowledge, skills, and abilities necessary to perform the corresponding job roles. Additionally, we highly recommend that you use the certification practice exam to prepare for the Security+ SY0-701 exam.

The questions on certification exams are protected to maintain the integrity of the exam. While the practice exam does not include the exact questions, it will help measure your understanding of the topics covered in the CompTIA objectives. You should review the Security+ SY0-701 exam objectives and make sure you are comfortable with each topic and objective listed. After you take the practice exam, the objectives for the exam are listed on the Report screen. You can use the report to focus your studies to prepare for the exam.

Is there a student discount for the Security+ SY0-701 exam?

CompTIA does offer discounts on its exam vouchers. Also, academic pricing for students is available through the academic store. Check the exam information page and www.comptia.org/blog/voucher-discount for more information.

TestOut is pleased to offer students a 10% off coupon code for exam vouchers purchased from CompTIA's online marketplace.

Is the Security+ SY0-701 exam an adaptive exam?

The Security+ SY0-701 exam and the certification practice exam are not adaptive tests. The certification practice exam is the best way to prepare for taking the certification exam. The content of the objectives for an exam is more comprehensive than any single adaptive test. Adaptive tests are too short to give you a thorough review and the chance to practice taking the test. You need to understand all of the questions before you take the certification exam.

An adaptive exam begins by giving you an easy-to-moderate question. If you answer the question correctly, it gives you a more difficult question. With each correct answer, the difficulty of the next question increases. On the other hand, if you answer the second question incorrectly, the next questions will be easier. The test changes the question difficulty until it determines your skill level.

There are two primary characteristics you will notice as you take an adaptive exam:

- You cannot skip questions or review previously answered questions. This means that you need to take more time to answer each question carefully before going on to the next question. Adaptive exams display a warning screen at the beginning of the exam stating that you will not be allowed to review previous questions.

- Adaptive tests are typically shorter than traditional exams. The current adaptive exams range between 15 and 35 questions.

How is the exam administered?

Certification exams are all computer-based. At the beginning of the exam, you will have an opportunity to view a tutorial on the exam software. Time spent reviewing the tutorial does not count towards the time you have to take the exam.

What types of questions are on the exam?

The Security+ SY0-701 exam has mostly multiple-choice questions. Additionally, there are a few performance-based questions at the start of the exam and some drag-and-drop activities. The CompTIA website provides sample questions and information on performance-based questions.

How many questions are on an exam, and how long do I have to complete it?

The Security+ SY0-701 exam has a maximum of 90 questions and 90 minutes to complete the exam.

How soon after I take the exam will I know whether I have passed it?

You will receive results as soon as you have completed the exam. The testing program provides immediate feedback and automatically generates a report showing the required passing score and your score. Pick up your exam report before you leave the testing center. You should keep this report in case there are any discrepancies in your certification program.

If I do not pass the exam, can I retake it?

For the Security+ SY0-701 exam, see the CompTIA website for exam policies. Typically, there is no waiting period before attempting to retake the exam a second time; however, the 3rd and subsequent attempts have a waiting period. See the CompTIA website for specific information on the retake policy.

A.1.5 Hints and Tips for taking the Security+ SY0-701 Exam

Follow these tips to make your exam experience less stressful and more successful:

When	Tips
Before the Exam	Before you take the exam, try these tips: Prepare a short review sheet for the exam. It should contain reference tables and information that you have trouble remembering. Shortly before you start the exam, study your notes as a last-minute review. Arrive 20 minutes early and relax for a few minutes before the exam. Take a deep breath. Look at the review sheet one last time. You will make fewer mistakes if you are not tense and rushed. Before the exam starts, review the exam tutorial to familiarize yourself with the exam. The time you spend on the orientation exam does not count toward the test time. If you have any questions, ask the exam administrator before the exam begins. The exam is timed, so make sure you ask questions before the test begins.
During the Exam	Once the exam has started: If you are unsure of a question's answer, eliminate the obviously incorrect answers first. Eliminating the obvious makes it easier for you to try to select the correct answer, especially if you have to guess. If you do not know, guess! Be sure you answer all of the questions before you finish. Unanswered questions are wrong and scored as incorrect answers. If you are unsure of an answer, make an educated guess. There is no extra penalty for incorrect answers. If you have time, review your answers before going on to the next question. A word of caution: be absolutely sure before you change an answer! If you are positive that your answer is wrong, change it. But if you are not sure and cannot explain to yourself why you need to change an answer, leave it. Most of the time, your first instinct is correct.
Use Scratch Pad	You will be given an erasable marker and a whiteboard or laminated paper to use during the exam. Follow these tips for using the scratch pad: Immediately after the exam starts, write down anything that could be a useful reference during the exam. This is the time to remember what you studied on your review sheet.

	The information on the review sheet should be fresh in your mind, because you just did a quick review. Write lists, reference tables, and any other vital information on the paper. Do not spend a lot of time, just a minute or two writing down reference material. The list of information will save you time as you answer the questions. While answering questions, use the scratch pad to draw diagrams. A question may be easier to answer after you see a diagram.
Retake an Exam	If you do not pass the exam: Use the score report on your transcript to identify the areas to focus further study. Think carefully about the exam and make notes about the questions that you could not answer. Do this as soon as possible after taking the exam. Look up the correct answers in your study materials. You may get the same or similar questions the next time. Do not wait too long to retake the exam. You already know much of the material, and you may forget what you know if you wait too long.

A.2 CompTIA Security+ Domain Review (20 Questions)

This section contains five domain practice reviews related to the CompTIA Security+ SY0-701 domains.

The domain practices:

- Have 20 questions per domain review
- Are randomly generated
- Have no time limit
- Are not pass/fail
- Have a percentage score at the end of the session

This section is meant to be used as a practice review only, not as a graded assignment. If a graded assignment is needed, use the CompTIA Security+ SY0-701 Certification Practice Exam score.

With these study questions, you can:

- Check your answers as you go to obtain immediate feedback.
- Skip questions, return to previous questions, and mark questions for later review.
- Repeat questions as many times as needed.

After you finish the study questions, you can:

- Review each question by selecting the Individual Responses option.

Expand questions marked as incorrect to see the correct answer and an accompanying explanation.

Print the score report to use for further review.

A.2.1 Security+ SY0-701 Domain 1: General Security Concepts (Section Quiz)

A.2.2 Security+ SY0-701 Domain 2: Threats, Vulnerabilities, and Mitigations (Section Quiz)

A.2.3 Security+ SY0-701 Domain 3: Security Architecture (Section Quiz)

A.2.4 Security+ SY0-701 Domain 4: Security Operations (Section Quiz)

A.2.5 Security+ SY0-701 Domain 5: Security Program Management and Oversight (Section Quiz)

A.3 CompTIA Security+ Domain Review (All Questions)

This section contains all the TestOut practice questions related to the Security+ SY0-701 exam. Use these questions to prepare yourself for the Security+ SY0-701 exam.

The questions in this section are not randomly generated. You will receive the same set of questions each time you open a new practice session.

This practice review has:

1374 questions

No time limit

No pass/fail

A percentage score at the end of the session

This section is meant to be used as a practice review only, not as a graded assignment. If a graded assignment is needed, use the Security+ SY0-701 Certification Practice Exam score.

With these study questions, you can:

Check your answers as you go to obtain immediate feedback.

Skip questions, return to previous questions, and mark questions for later review.

Repeat questions as many times as needed.

After you finish the study questions, you can:

Review each question by selecting the Individual Responses option.

Expand questions marked as incorrect to see the correct answer and an explanation of the correct answer.

Print the score report to use for further review.

A.3.1 Security+ SY0-701 Domain 1: General Security Concepts (Section Quiz)

A.3.2 Security+ SY0-701 Domain 2: Threats, Vulnerabilities, and Mitigations (Section Quiz)

A.3.3 Security+ SY0-701 Domain 3: Security Architecture (Section Quiz)

A.3.4 Security+ SY0-701 Domain 4: Security Operations (Section Quiz)

A.3.5 Security+ SY0-701 Domain 5: Security Program Management and Oversight (Section Quiz)

A.4 CompTIA Security+ SY0-701 Certification Practice Exam (Section Quiz)

B.0 TestOut Security Pro - Practice Exams

B.1 Prepare for TestOut Security Pro Certification

It is important to prepare for an exam by studying the course material, practicing skills, and committing new concepts to memory. You can use the instructions and tests in this course to help you prepare more efficiently.

We recommend that you take the following steps as you prepare for the TestOut Security Pro Certification exam:

Step	Description
Study the course material	The course materials include text lessons, demonstrations, video lessons, and hands-on labs. As you work through the course, follow these hints for the effective study: Review the learning and exam objectives on each section page. The objectives outline the knowledge and skills you will need for the official certification exam. Watch the videos. Watch the demonstrations. Read all text lesson fact pages. Practice the tasks in the lab simulations until you feel comfortable with your ability to complete them. Avoid skipping any sections unless you can easily pass the Practice Questions at the end of each section. Even if you already know the material, a review can always be helpful when preparing for an exam.
Review the certification exam domains and objectives	Review the domains and objectives for the TestOut Security Pro Certification provided in this section.
Take the domain practice exams	The domain practice exams group the performance-based labs by domain and help assess your understanding of a particular TestOut Security Pro Certification domain and the corresponding objectives.
Take the certification practice exam	After you are confident with your ability to complete the labs, take the certification practice exam to assess your preparedness to take the certification exam. This exam has roughly the same number of questions and time limit as the TestOut Security Pro Certification. Practice questions are designed to assess your knowledge as it relates to the exam objectives.

	Based on your practice exam results, review the course material for questions that you missed. Focus your time on understanding the topics covered in the objectives and not on memorizing answers, as the actual certification exam will have a different set of questions. When your practice exam scores are consistently over 95%, and you feel confident in your understanding of the exam objectives and topics, the next step is to take the certification exam.
Schedule and take the certification exam	The TestOut Security Pro Certification is scheduled through LabSim. If you are taking this course through an instructor, contact your instructor to schedule your exam.

B.1.1 Pro Exam Objectives

The Security Pro course and certification exam cover the following TestOut Security Pro objectives:

#	Domain	Module.Section
1.0	**Identity Management and Authentication**	
1.1	Manage identity 1.1.1 Manage Windows local and domain users and groups 1.1.2 Manage Linux users and groups 1.1.3 Manage Active Directory OUs	4.4, 4.6, 4.7 8.1
1.2	Harden authentication 1.2.1 Configure account policies 1.2.2 Manage account password 1.2.3 Secure default and local accounts 1.2.4 Enforce User Account Control (UAC) 1.2.5 Configure and link Group Policy Objects (GPO)	4.4, 4.5, 4.6 8.1 12.3
2.0	**Physical and Network Security**	
2.1	Harden physical access 2.1.1 Implement physical security 2.1.2 Install and configure a security appliance	5.2, 5.3, 5.4 6.1 8.1, 8.3

	2.1.3 Install and configure a firewall 2.1.4 Create and configure a screened subnet 2.1.5 Configure Network Address Translation (NAT)	
2.2	Harden network devices 2.2.1 Configure and access a switch 2.2.2 Configure and access a wireless network 2.2.3 Configure and access a Virtual Private Network (VPN) 2.2.4 Harden a wireless network 2.2.5 Configure router security 2.2.6 Bring Your Own Device (BYOD) security 2.2.7 Create and connect to a Virtual Local Area Network (VLAN)	4.8 5.5, 5.7, 5.9, 5.10 6.2 8.4, 8.5, 8.6 10.5, 10.6, 10.7
3.0	**Host and Application Defense**	
3.1	Harden computer systems 3.1.1 Configure file system inheritance 3.1.2 Configure anti-virus protection 3.1.3 Configure NTFS permissions 3.1.4 Configure Windows Update	2.3 7.1, 7.3 8.1, 8.2
3.2	Implement application defenses 3.2.1 Implement an application allow list 3.2.2 Implement Data Execution Prevention (DEP) 3.2.3 Configure web application security 3.2.4 Configure email filters and settings 3.2.5 Configure browser settings	8.1, 8.7, 8.8, 8.9 10.6, 10.9
3.3	Implement virtualization 3.3.1 Create virtual machines 3.3.2 Create virtual switches	10.1, 10.2
4.0	**Data Security**	

4.1	Protect and Maintain Data files 4.1.1 Perform data backups and recovery 4.1.2 Implement redundancy	9.4, 9.5
4.2	Implement Encryption Technologies 4.2.1 Encrypt data communications 4.2.2 Encrypt files 4.2.3 Manage certificates	3.1, 3.2, 3.3, 3.4, 3.5 7.3 8.2, 8.7
5.0	**Audit and Security Assessment**	
5.1	Implement logging and auditing 5.1.1 Configure advanced audit policy 5.1.2 Enable device logs	7.3 9.2 12.3
5.2	Assessment techniques 5.2.1 Implement intrusion detection 5.2.2 Identify social engineering 5.2.3 Scan for vulnerabilities 5.2.4 Analyze network attacks 5.2.5 Analyze password attacks	2.2 6.3, 6.5, 6.6 7.1, 7.2, 7.4 9.1, 9.2 10.1

B.1.2 Pro Exam Objectives by Course Section

The Security Pro course covers the following TestOut Security Pro exam objectives:

Section	Title	Objectives
1.0	**Security Concepts**	
1.1	Security Introduction	
1.2	Security Controls	
1.3	Use the Simulator	
2.0	**Threats, Vulnerabilities, and Mitigations**	

2.1	Understanding Attacks	
2.2	Social Engineering	5.2 Assessment techniques • 5.2.2 Identify social engineering
2.3	Malware	3.1 Harden computer systems • 3.1.2 Configure anti-virus protection
3.0	**Cryptographic Solutions**	
3.1	Cryptography	4.2 Implement Encryption Technologies • 4.2.1 Encrypt data communications
3.2	Cryptography Implementations	4.2 Implement Encryption Technologies • 4.2.1 Encrypt data communications • 4.2.2 Encrypt files
3.3	Hashing	4.2 Implement Encryption Technologies • 4.2.1 Encrypt data communications
3.4	Encryption	4.2 Implement Encryption Technologies • 4.2.2 Encrypt files
3.5	Public Key Infrastructure	4.2 Implement Encryption Technologies • 4.2.3 Manage certificates
4.0	**Identity and Access Management**	
4.1	Access Control Models	
4.2	Authentication	
4.3	Authorization	
4.4	Active Directory Overview	1.1 Manage identity • 1.1.1 Manage Windows local and domain users and groups • 1.1.3 Manage Active Directory OUs

		1.2 Harden authentication • 1.2.5 Configure and link Group Policy Objects (GPO)
4.5	Hardening Authentication	1.2 Harden authentication • 1.2.1 Configure account policies • 1.2.2 Manage account password • 1.2.3 Secure default and local accounts • 1.2.4 Enforce User Account Control (UAC) • 1.2.5 Configure and link Group Policy Objects (GPO)
4.6	Linux Users	1.1 Manage identity • 1.1.2 Manage Linux users and groups 1.2 Harden authentication • 1.2.2 Manage account password • 1.2.3 Secure default and local accounts
4.7	Linux Groups	1.1 Manage identity • 1.1.2 Manage Linux users and groups
4.8	Remote Access	2.2 Harden network devices • 2.2.3 Configure and access a Virtual Private Network (VPN)
4.9	Network Authentication	
5.0	**Network Architecture**	
5.1	Enterprise Network Architecture	
5.2	Security Appliances	2.1 Harden physical access • 2.1.2 Install and configure a security appliance • 2.1.4 Create and configure a screened subnet
5.3	Screened Subnets	2.1 Harden physical access • 2.1.4 Create and configure a screened subnet

5.4	Firewalls	2.1 Harden physical access • 2.1.3 Install and configure a firewall
5.5	Virtual Private Networks	2.2 Harden network devices • 2.2.3 Configure and access a Virtual Private Network (VPN) • 2.2.4 Harden a wireless network
5.6	Network Access Control	
5.7	Network Device Vulnerabilities	2.2 Harden network devices • 2.2.1 Configure and access a switch
5.8	Network Applications	
5.9	Switch Security and Attacks	2.2 Harden network devices • 2.2.1 Configure and access a switch
5.10	Router Security	2.2 Harden network devices • 2.2.5 Configure router security
6.0	**Resiliency and Site Security**	
6.1	Physical Threats	2.1 Harden physical access • 2.1.1 Implement physical security
6.2	Monitoring and Reconnaissance	2.2 Harden network devices • 2.2.4 Harden a wireless network
6.3	Intrusion Detection	5.2 Assessment techniques • 5.2.1 Implement intrusion detection
6.4	Protocol Analyzers	
6.5	Analyzing Network Attacks	5.2 Assessment techniques • 5.2.4 Analyze network attacks • 5.2.5 Analyze password attacks
6.6	Analyzing Password Attacks	5.2 Assessment techniques • 5.2.2 Identify social engineering • 5.2.5 Analyze password attacks
7.0	**Vulnerability Management**	

7.1	Vulnerability Management	3.1 Harden computer systems • 3.1.4 Configure Windows Update 5.2 Assessment techniques • 5.2.3 Scan for vulnerabilities
7.2	Vulnerability Scanning	5.2 Assessment techniques • 5.2.3 Scan for vulnerabilities
7.3	Alerting and Monitoring	3.1 Harden computer systems • 3.1.2 Configure anti-virus protection 4.2 Implement Encryption Technologies • 4.2.1 Encrypt data communications • 4.2.2 Encrypt files 5.1 Implement logging and auditing • 5.1.2 Enable device logs
7.4	Penetration Testing	5.2 Assessment techniques • 5.2.3 Scan for vulnerabilities
8.0	**Network and Endpoint Security**	
8.1	Operating System Hardening	1.1 Manage identity • 1.1.1 Manage Windows local and domain users and groups

		1.2 Harden authentication • 1.2.5 Configure and link Group Policy Objects (GPO) 2.1 Harden physical access • 2.1.3 Install and configure a firewall 3.1 Harden computer systems • 3.1.2 Configure anti-virus protection • 3.1.4 Configure Windows Update 3.2 Implement application defenses • 3.2.

		1 Implement an application allow list
8.2	File Server Security	3.1 Harden computer systems • 3.1.1 Configure file system inheritance

- 3.1.3 Configure NTFS permissions
- 4.2 Implement Encryption Techn

		• 4.2.2 Encrypt files ologies
8 3	Linux Host Security	2.1 Harden physical access

	Wireless Controller	

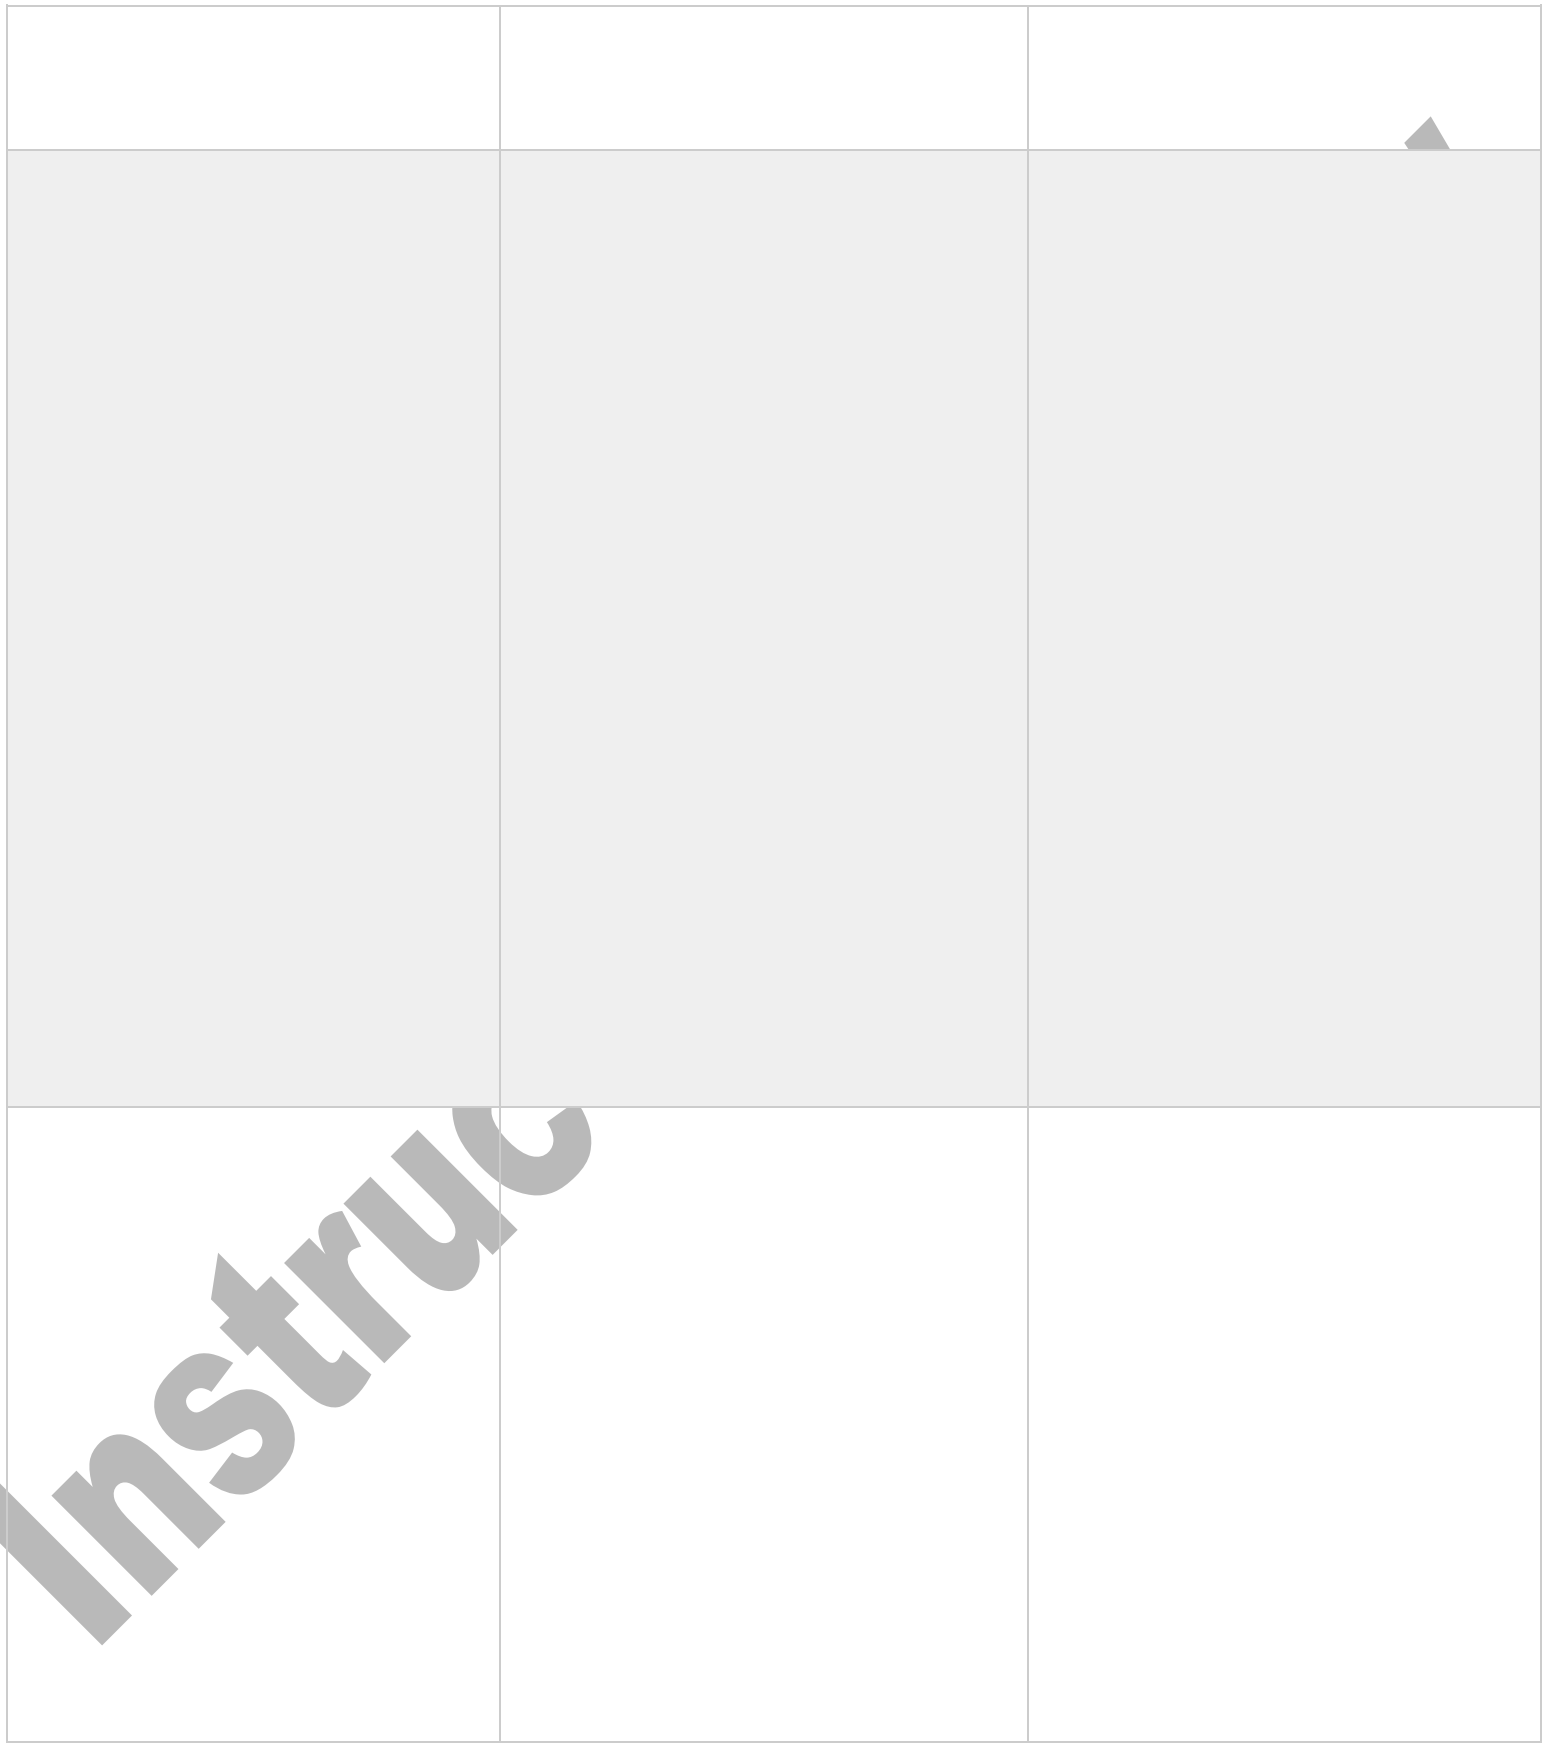

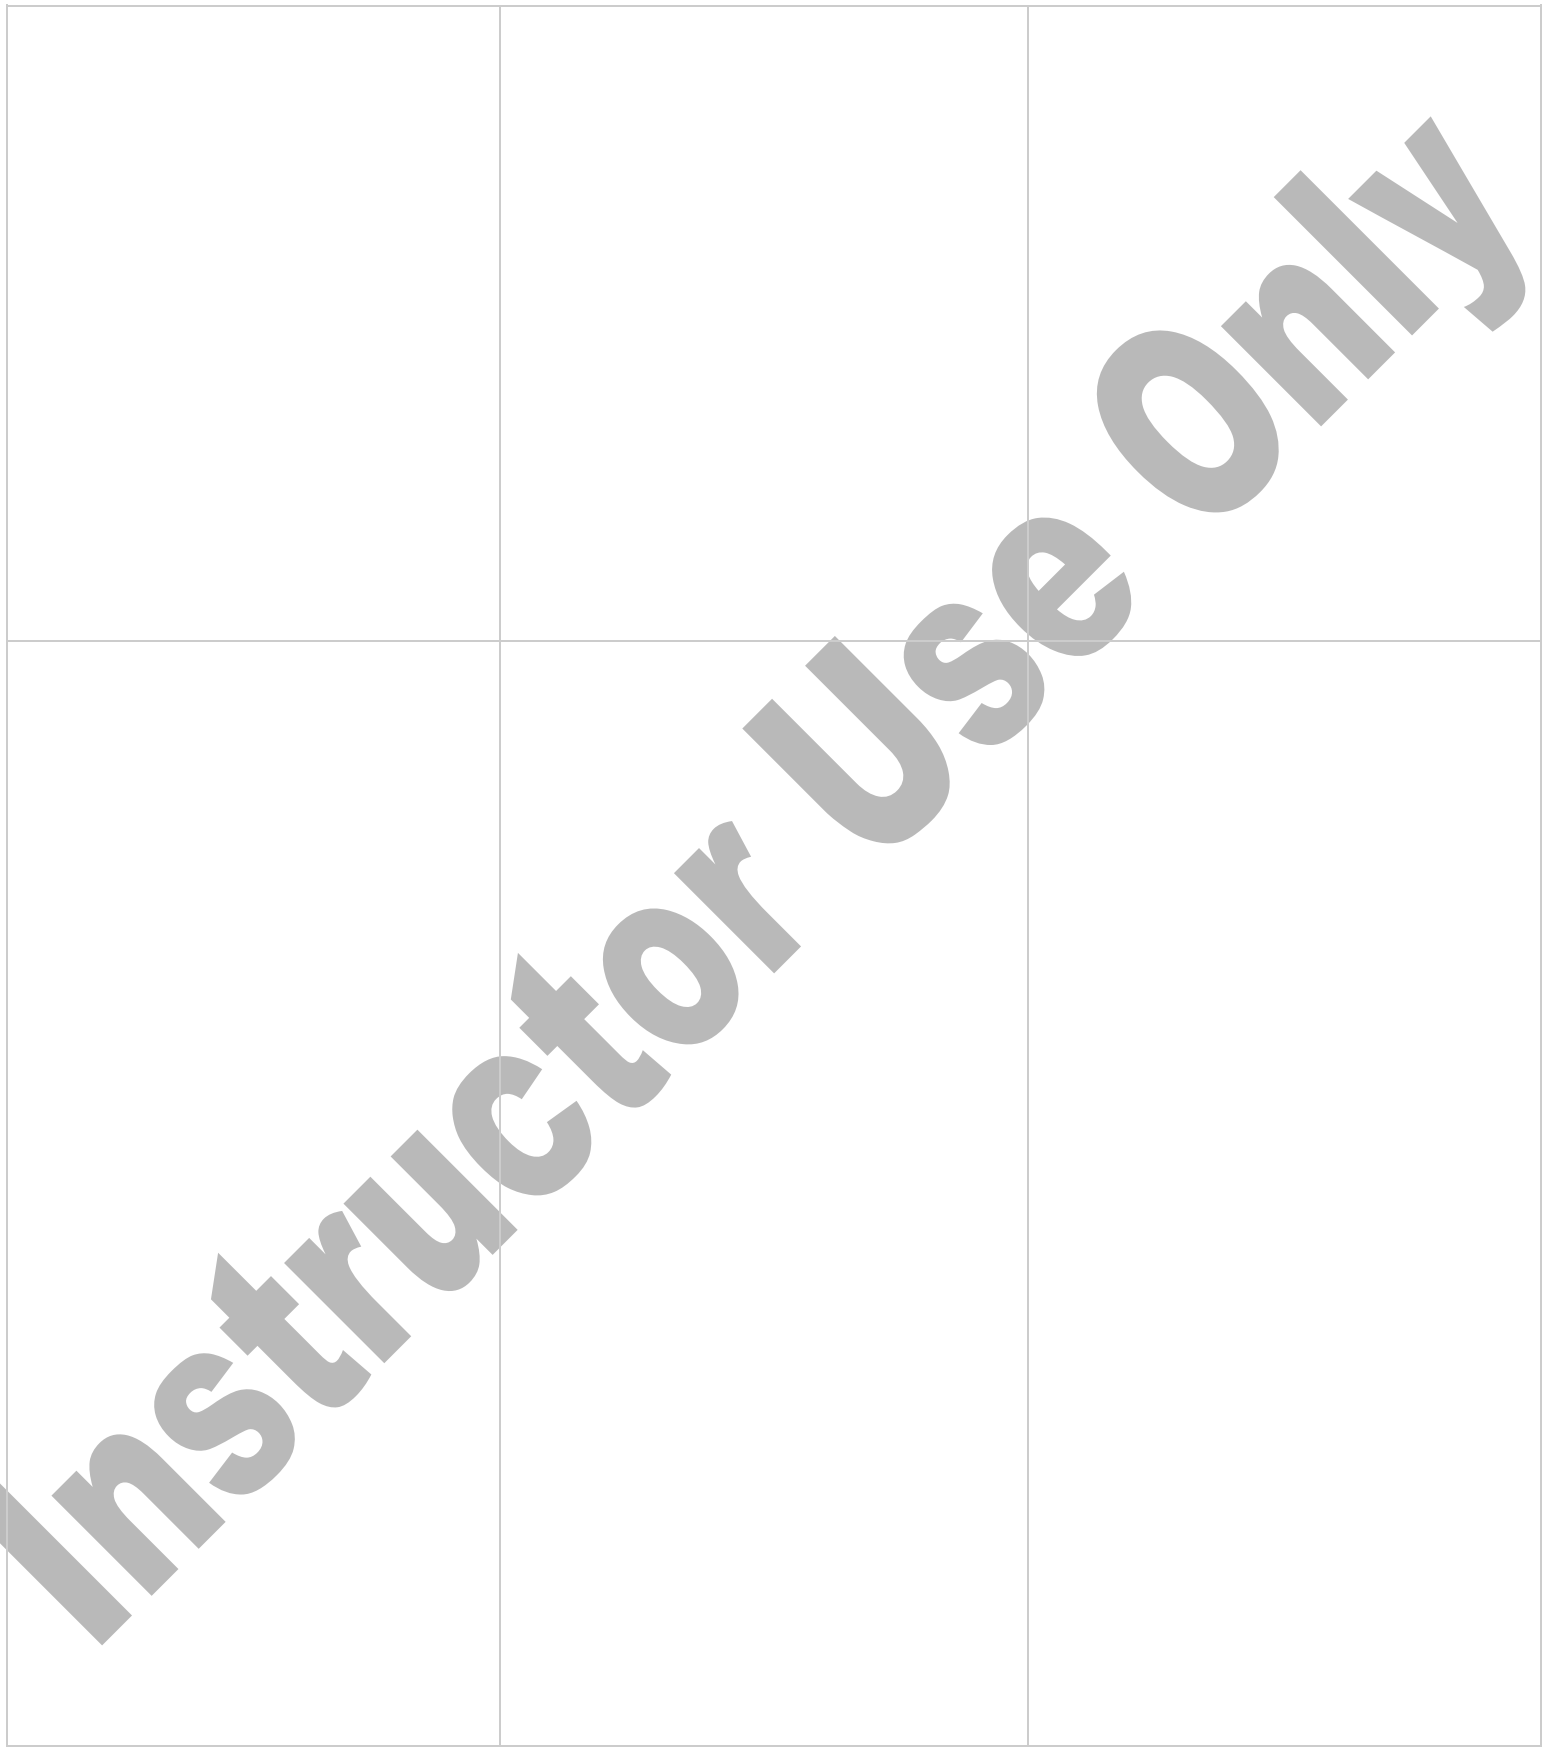

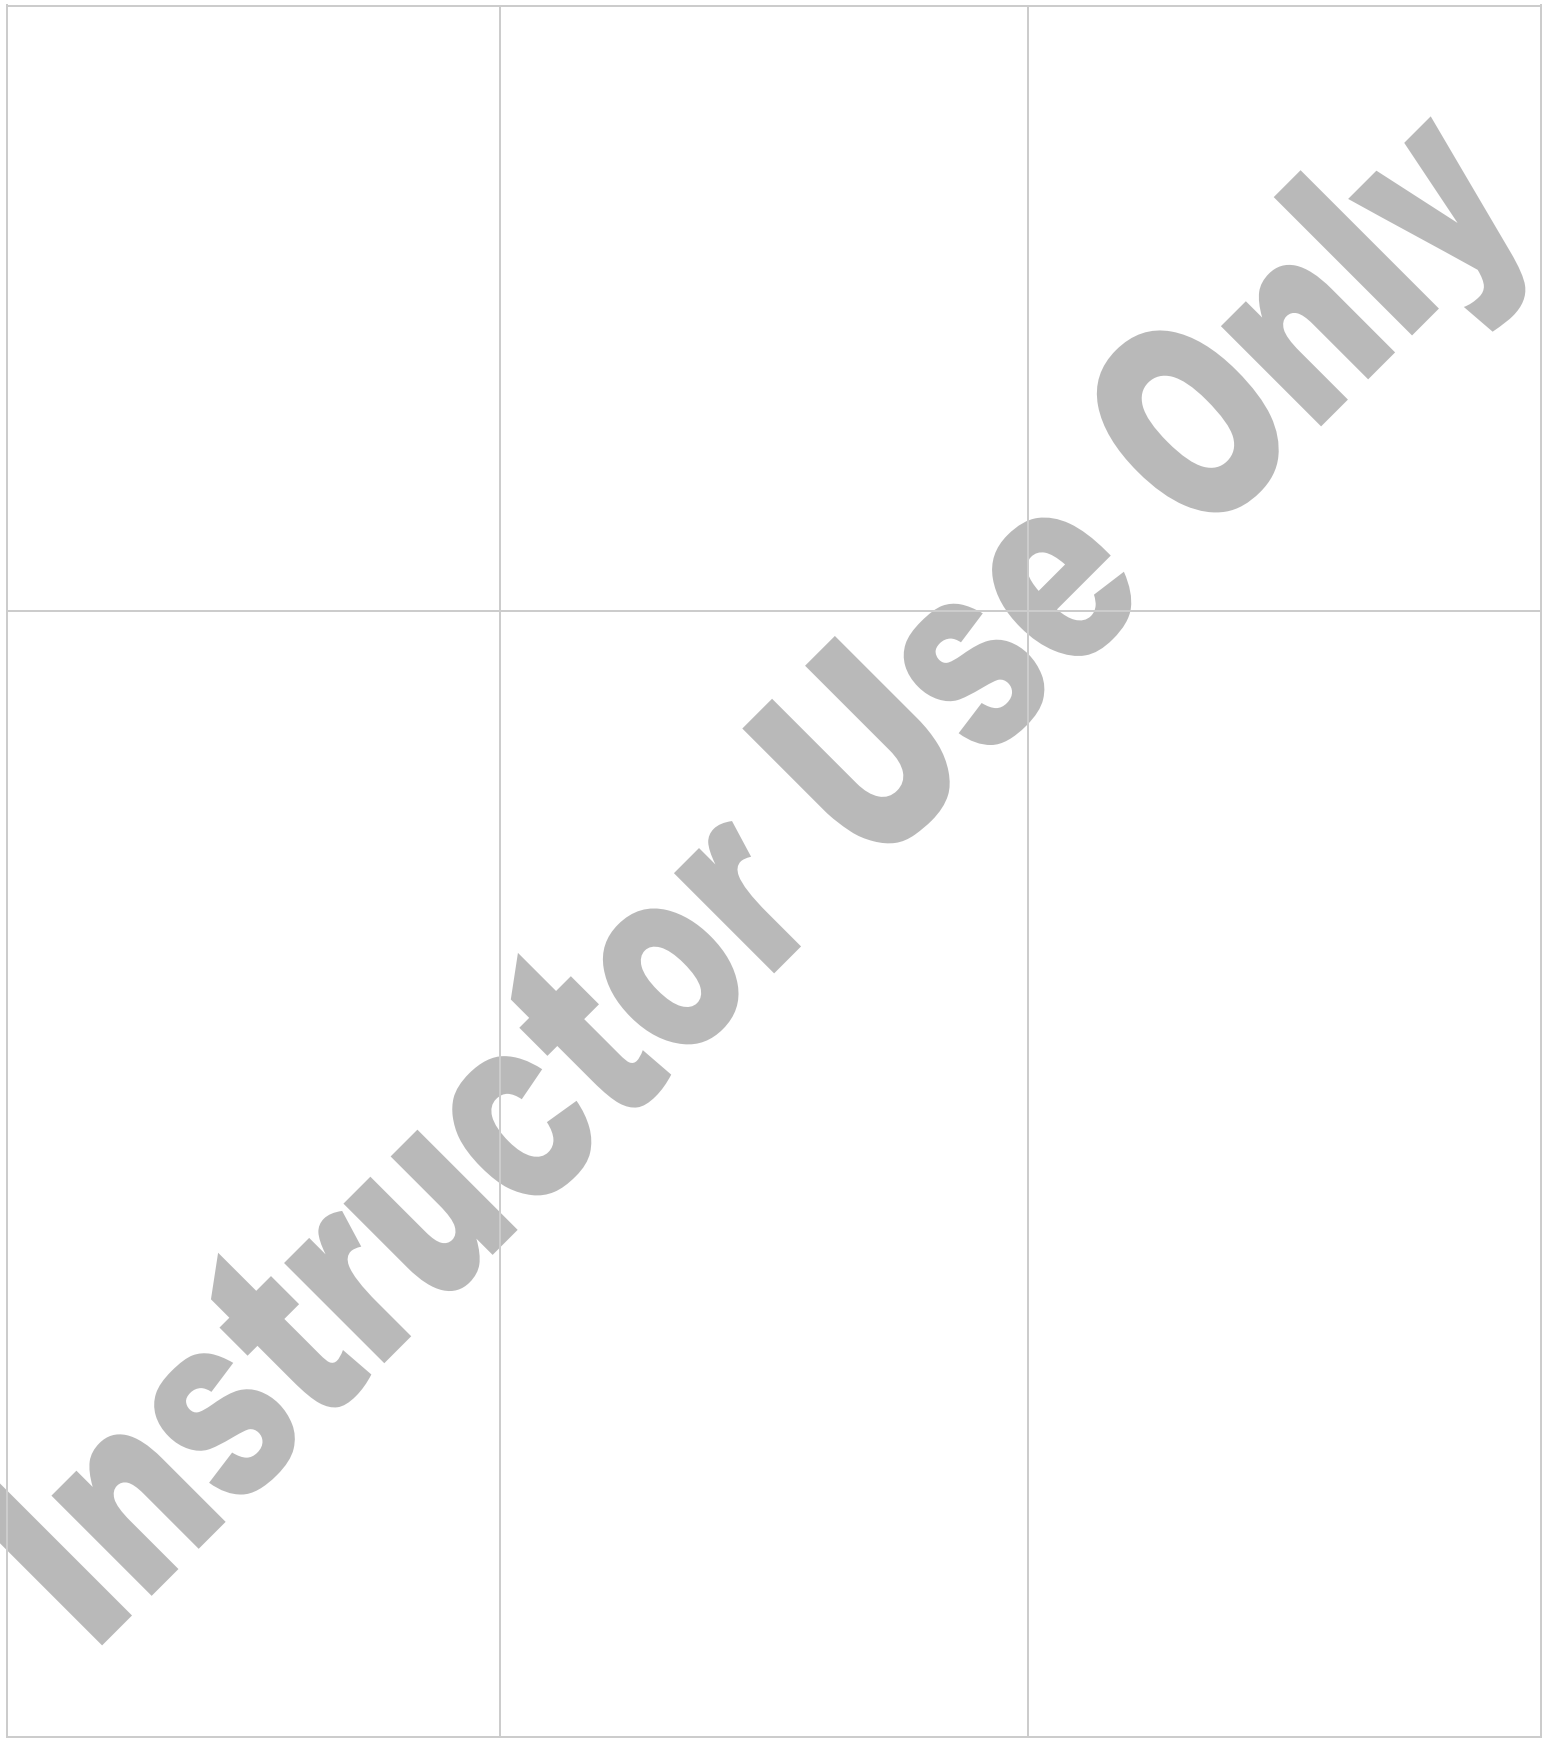

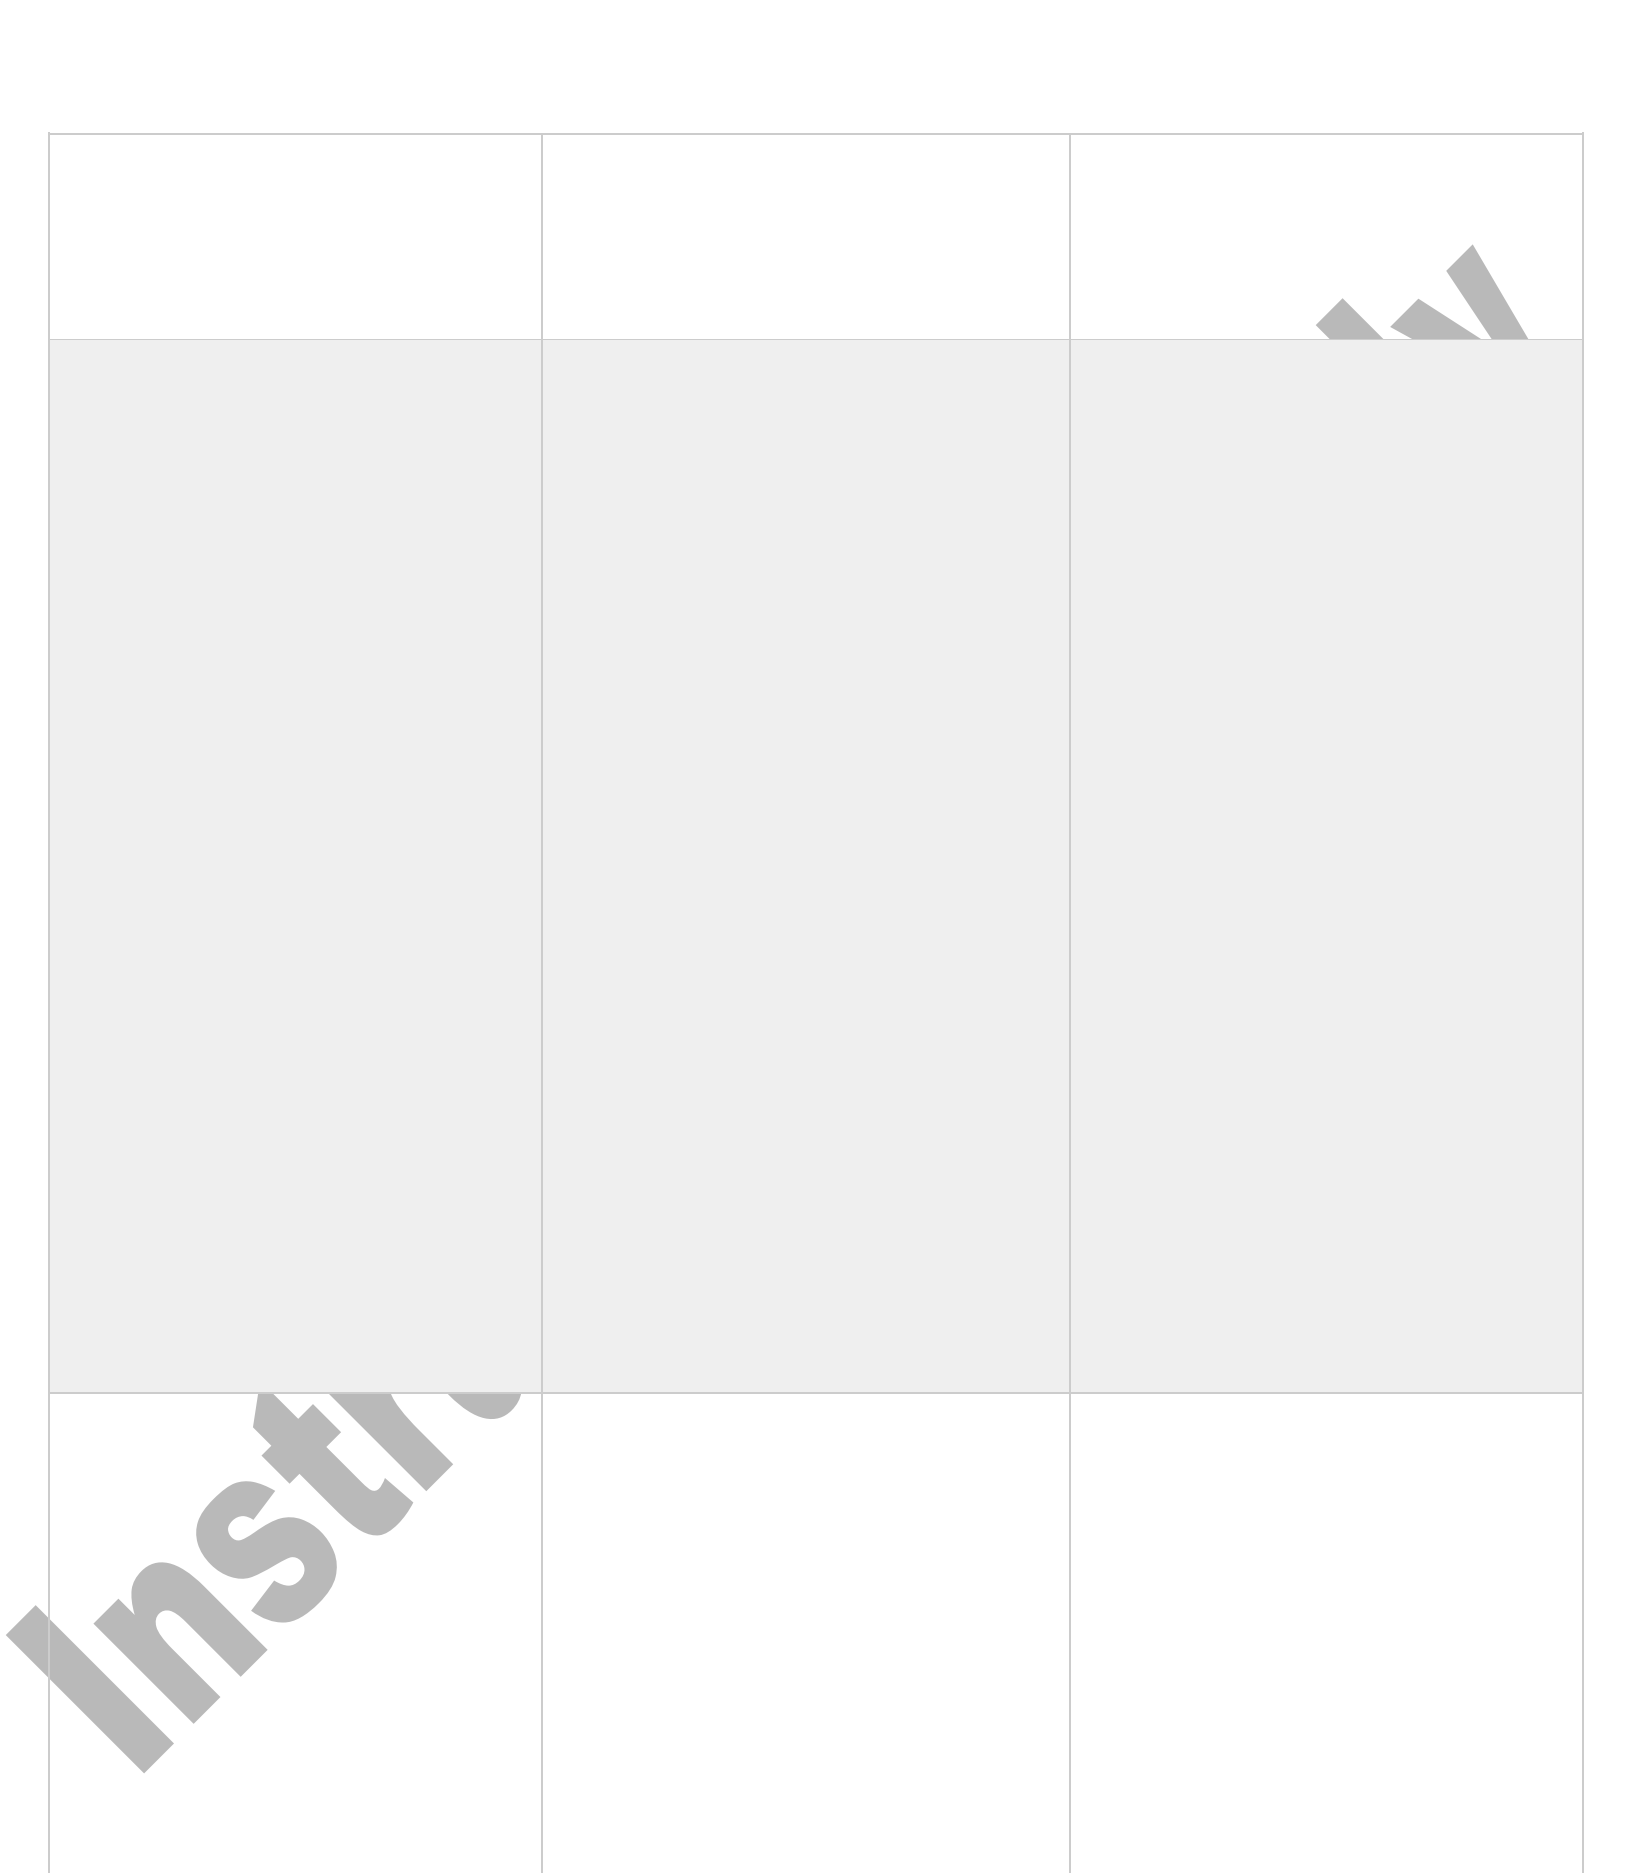

B.1.3 How to take the Pro Exam

To take the TestOut Security Pro Certification exam, you must first schedule the exam with an instructor or an administrator.

You are encouraged to budget your time and to keep moving through the exam so you can complete it within the time limit. Begin an exam at the scheduled exam time. Follow these steps to start the exam:

- Log in to LabSim.
- Select the **Certifications** tab from the top navigation menu.
- Navigate to the certification.
- Select **Start Exam** and read the instructions.
- Select **Start Exam** when you are ready to begin the exam.

Exam Notes

- When you purchase a TestOut course, the Pro Certification exam is included with the course. The course license must be active for you to take the exam.
- Unlike the practice exams, the TestOut Pro certification exams are assessment exams that do not allow students to check their answers or get instant feedback.
- The exam may be attempted only once per exam voucher.
- If you do not pass the exam, you may purchase a retake exam voucher online.
- Do not click **Start Exam** until you are ready to complete the certification exam. Starting the exam and exiting will use the exam voucher.

B.1.4 Pro Exam FAQs

TestOut Pro Certification Frequently Asked Questions (FAQs)

For the most up-to-date TestOut Pro Certification FAQs, visit TestOut Pro Certification at https://w3.testout.com/certification/pro-exams/resources/pro-exam-faqs .

B.2 TestOut Security Pro Domain Review

This section is a practice review that contains all the questions related to the TestOut Security Pro exam. Use these questions to prepare for the TestOut Security Pro exam.

The questions in this section are not randomly generated. You will receive the same set of questions each time you open a new practice session.

This practice review has:

- 84 lab simulations
- No time limit
- No pass/fail score
- A percentage score at the end of the session

The student's score should not be used as part of the student's grade. This section is meant to be used as a practice review only. If you need to include a score in the student's grade, use the TestOut Security Pro Certification Practice Exam score.

After you finish the practice review, you can print the score report to use for further review.

B.2.1 Pro Domain 1: Identity Management and Authentication (Section Quiz)

B.2.2 Pro Domain 2: Physical and Network Security (Section Quiz)

B.2.3 Pro Domain 3: Host and Application Defense (Section Quiz)

B.2.4 Pro Domain 4: Data Security (Section Quiz)

B.2.5 Pro Domain 5: Audit and Security Assessment (Section Quiz)

B.3 TestOut Security Pro Certification Practice Exam (Section Quiz)