Chapter 1

Today's Cybersecurity Analyst

THE COMPTIA CYBERSECURITY ANALYST (CYSA+) EXAM OBJECTIVES COVERED IN THIS CHAPTER INCLUDE:

✓ **Domain 1.0: Security Operations**

- 1.5 Explain the importance of efficiency and process improvement in security operations
  - Standardize processes
  - Streamline operations
  - Technology and tool integration
  - Single pane of glass

✓ **Domain 2.0: Vulnerability Management**

- 2.1 Given a scenario, implement vulnerability scanning methods and concepts
  - Static vs. dynamic (reverse engineering)

Cybersecurity analysts are responsible for protecting the confidentiality, integrity, and availability of information and information systems used by their organizations. Fulfilling this responsibility requires a commitment to a defense-in-depth approach to information security that uses multiple, overlapping security controls to achieve each cybersecurity objective. It also requires that analysts have a strong understanding of the threat environment facing their organization in order to develop a set of controls capable of rising to the occasion and answering those threats.

In the first section of this chapter, you will learn how to assess the cybersecurity threats facing your organization and determine the risk that they pose to the confidentiality, integrity, and availability of your operations. In the sections that follow, you will learn about controls that you can put in place to secure networks and endpoints and evaluate the effectiveness of those controls over time.

Cybersecurity Objectives

When most people think of cybersecurity, they imagine hackers trying to break into an organization's system and steal sensitive information, ranging from Social Security numbers and credit cards to top-secret military information. Although protecting sensitive information from unauthorized disclosure is certainly one element of a cybersecurity program, it is important to understand that cybersecurity actually has three complementary objectives, as shown in Figure 1.1.

FIGURE 1.1 The three key objectives of cybersecurity programs are confidentiality, integrity, and availability.

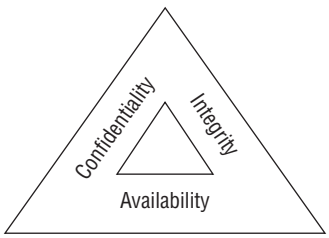

Confidentiality ensures that unauthorized individuals are not able to gain access to sensitive information. Cybersecurity professionals develop and implement security controls, including firewalls, access control lists, and encryption, to prevent unauthorized access to information. Attackers may seek to undermine confidentiality controls to achieve one of their goals: the unauthorized disclosure of sensitive information.

Integrity ensures that there are no unauthorized modifications to information or systems, either intentionally or unintentionally. Integrity controls, such as hashing and integrity monitoring solutions, seek to enforce this requirement. Integrity threats may come from attackers seeking the alteration of information without authorization or nonmalicious sources, such as a power spike causing the corruption of information.

Availability ensures that information and systems are ready to meet the needs of legitimate users at the time those users request them. Availability controls, such as fault tolerance, clustering, and backups, seek to ensure that legitimate users may gain access as needed. Similar to integrity threats, availability threats may come either from attackers seeking the disruption of access or nonmalicious sources, such as a fire destroying a datacenter that contains valuable information or services.

Cybersecurity analysts often refer to these three goals, known as the CIA Triad, when performing their work. They often characterize risks, attacks, and security controls as meeting one or more of the three CIA Triad goals when describing them.

Privacy vs. Security

Privacy and security are closely related concepts. We just discussed the three major components of security: confidentiality, integrity, and availability. These goals are all focused on the ways that an organization can protect its own data. Confidentiality protects data from unauthorized disclosure. Integrity protects data from unauthorized modification. Availability protects data from unauthorized denial of access.

Privacy controls have a different focus. Instead of focusing on ways that an organization can protect its own information, privacy focuses on the ways that an organization can use and share information that it has collected about individuals. This data, known as *personally identifiable information* (*PII*), is often protected by regulatory standards and is always governed by ethical considerations. Organizations seek to protect the security of private information and may do so using the same security controls that they use to protect other categories of sensitive information, but privacy obligations extend beyond just security. Privacy extends to include the ways that an organization uses and shares the information that it collects and maintains with others.

Exam Note

Remember that privacy and security are complementary and overlapping, but they have different objectives. This is an important concept on the exam.

The *Generally Accepted Privacy Principles (GAPP)* outline 10 privacy practices that organizations should strive to follow:

- **Management** says that the organization should document its privacy practices in a privacy policy and related documents.
- **Notice** says that the organization should notify individuals about its privacy practices and inform individuals of the type of information that it collects and how that information is used.
- **Choice and consent** says that the organization should obtain the direct consent of individuals for the storage, use, and sharing of PII.
- **Collection** says that the organization should collect PII only for the purposes identified in the notice and consented to by the individual.
- **Use, retention, and disposal** says that the organization should only use information for identified purposes and may not use information collected for one stated purpose for any other nondisclosed purpose.
- **Access** says that the organization should provide individuals with access to any information about that individual in the organization's records, at the individual's request.
- **Disclosure** says that the organization will disclose information to third parties only when consistent with notice and consent.
- **Security** says that PII will be protected against unauthorized access.
- **Quality** says that the organization will maintain accurate and complete information.
- **Monitoring and enforcement** says that the organization will put business processes in place to ensure that it remains compliant with its privacy policy.

The GAPP principles are strong best practices for building a privacy program. In some jurisdictions and industries, privacy laws require the implementation of several of these principles. For example, the European Union's (EU) General Data Protection Regulation (GDPR) requires that organizations handling the data of EU residents process personal information in a way that meets privacy requirements.

Evaluating Security Risks

Cybersecurity risk analysis is the cornerstone of any information security program. Analysts must take the time to thoroughly understand their own technology environments and the external threats that jeopardize their information security. A well-rounded cybersecurity risk

assessment combines information about internal and external factors to help analysts understand the threats facing their organization and then design an appropriate set of controls to meet those threats.

Before diving into the world of risk assessment, we must begin with a common vocabulary. You must know three important terms to communicate clearly with other risk analysts: vulnerabilities, threats, and risks.

A *vulnerability* is a weakness in a device, system, application, or process that might allow an attack to take place. Vulnerabilities are internal factors that may be controlled by cybersecurity professionals. For example, a web server that is running an outdated version of the Apache service may contain a vulnerability that would allow an attacker to conduct a denial-of-service (DoS) attack against the websites hosted on that server, jeopardizing their availability. Cybersecurity professionals within the organization have the ability to remediate this vulnerability by upgrading the Apache service to the most recent version that is not susceptible to the DoS attack.

A *threat* in the world of cybersecurity is an outside force that may exploit a vulnerability. For example, a hacker who would like to conduct a DoS attack against a website and knows about an Apache vulnerability poses a clear cybersecurity threat. Although many threats are malicious in nature, this is not necessarily the case. For example, an earthquake may also disrupt the availability of a website by damaging the datacenter containing the web servers. Earthquakes clearly do not have malicious intent. In most cases, cybersecurity professionals cannot do much to eliminate a threat. Hackers will hack and earthquakes will strike whether we like it or not.

A *risk* is the combination of a threat and a corresponding vulnerability. Both of these factors must be present before a situation poses a risk to the security of an organization. For example, if a hacker targets an organization's web server with a DoS attack but the server was patched so that it is not vulnerable to that attack, there is no risk because even though a threat is present (the hacker), there is no vulnerability. Similarly, a datacenter may be vulnerable to earthquakes because the walls are not built to withstand the extreme movements present during an earthquake, but it may be located in a region of the world where earthquakes do not occur. The datacenter may be vulnerable to earthquakes but there is little to no threat of earthquake in its location, so there is no risk.

The relationship between risks, threats, and vulnerabilities is an important one, and it is often represented by this equation:

$$\text{Risk} = \text{Threat} \times \text{Vulnerability}$$

This is not meant to be a literal equation where you would actually plug in values. Instead, it is meant to demonstrate the fact that risks exist only when there is both a threat and a corresponding vulnerability that the threat might exploit. If either the threat or vulnerability is zero, the risk is also zero. Figure 1.2 shows this in another way: risks are the intersection of threats and vulnerabilities.

FIGURE 1.2 Risks exist at the intersection of threats and vulnerabilities. If either the threat or vulnerability is missing, there is no risk.

Organizations should routinely conduct risk assessments to take stock of their existing risk landscape. The National Institute of Standards and Technology (NIST) publishes a guide for conducting risk assessments that is widely used throughout the cybersecurity field as a foundation for risk assessments. The document, designated NIST Special Publication (SP) 800-30, suggests the risk assessment process shown in Figure 1.3.

FIGURE 1.3 The NIST SP 800-30 risk assessment process suggests that an organization should identify threats and vulnerabilities and then use that information to determine the level of risk posed by the combination of those threats and vulnerabilities.

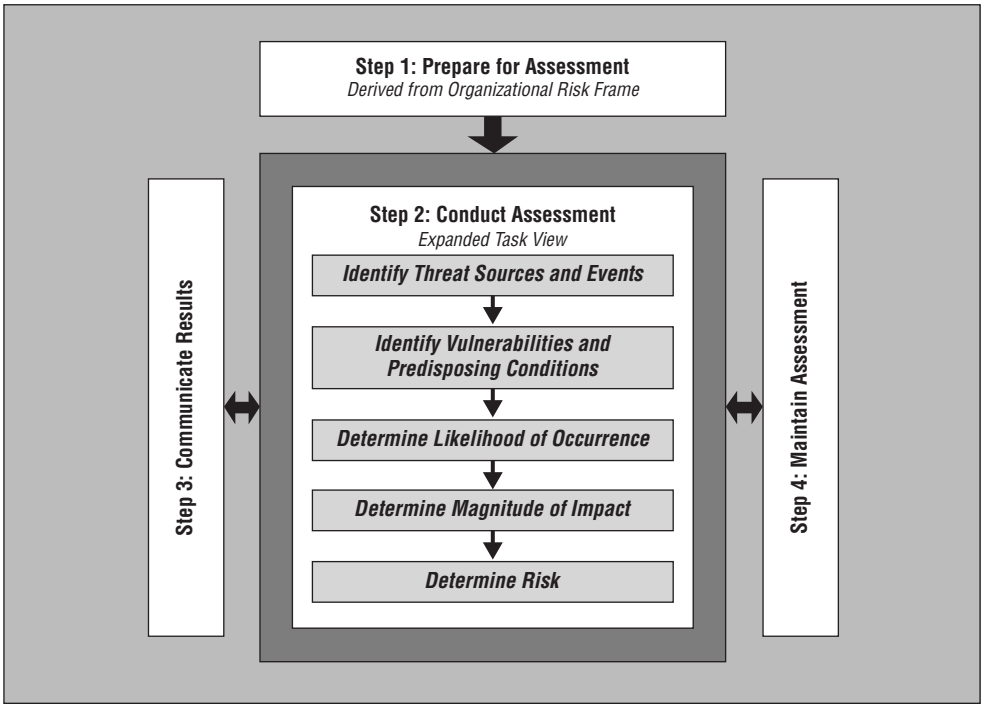

Source: NIST SP 800-30 / U.S Department of Commerce / Public Domain

Identify Threats

Organizations begin the risk assessment process by identifying the types of threats that exist in their threat environment. Although some threats, such as malware and spam, affect all organizations, other threats are targeted against specific types of organizations. For example, government-sponsored advanced persistent threat (APT) attackers typically target government agencies, military organizations, and companies that operate in related fields. It is unlikely that an APT attacker would target an elementary school.

NIST identifies four categories of threats that an organization might face and should consider in its threat identification process:

- **Adversarial threats** are individuals, groups, and organizations that are attempting to deliberately undermine the security of an organization. Adversaries may include trusted insiders, competitors, suppliers, customers, business partners, or even nation-states. When evaluating an adversarial threat, cybersecurity analysts should consider the capability of the threat actor to engage in attacks, the intent of the threat actor, and the likelihood that the threat will target the organization.
- **Accidental threats** occur when individuals doing their routine work mistakenly perform an action that undermines security. For example, a system administrator might accidentally delete a critical disk volume, causing a loss of availability. When evaluating an accidental threat, cybersecurity analysts should consider the possible range of effects that the threat might have on the organization.
- **Structural threats** occur when equipment, software, or environmental controls fail due to the exhaustion of resources (such as running out of gas), exceeding their operational capability (such as operating in extreme heat), or simply failing due to age. Structural threats may come from IT components (such as storage, servers, and network devices), environmental controls (such as power and cooling infrastructure), and software (such as operating systems and applications). When evaluating a structural threat, cybersecurity analysts should consider the possible range of effects that the threat might have on the organization.
- **Environmental threats** occur when natural or human-made disasters occur that are outside the control of the organization. These might include fires, flooding, severe storms, power failures, or widespread telecommunications disruptions. When evaluating environmental threats, cybersecurity analysts should consider common natural environmental threats to their geographic region, as well as how to appropriately prevent or counter human-made environmental threats.

The nature and scope of the threats in each of these categories will vary depending on the nature of the organization, the composition of its technology infrastructure, and many other situation-specific circumstances. That said, it may be helpful to obtain copies of the risk assessments performed by other, similar organizations as a starting point for an organization's own risk assessment or to use as a quality assessment check during various stages of the organization's assessment.

The Insider Threat

When performing a threat analysis, cybersecurity professionals must remember that threats come from both external and internal sources. In addition to the hackers, natural disasters, and other threats that begin outside the organization, rogue employees, disgruntled team members, and incompetent administrators also pose a significant threat to enterprise cybersecurity. As an organization designs controls, it must consider both internal and external threats.

NIST SP 800-30 provides a great deal of additional information to help organizations conduct risk assessments, including detailed tasks associated with each of these steps. This information is outside the scope of the Cybersecurity Analyst (CySA+) exam, but organizations preparing to conduct risk assessments should download and read the entire publication. It is available at `https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final`.

Identify Vulnerabilities

During the threat identification phase of a risk assessment, cybersecurity analysts focus on the external factors likely to impact an organization's security efforts. After completing threat identification, the focus of the assessment turns inward, identifying the vulnerabilities that those threats might exploit to compromise an organization's confidentiality, integrity, or availability.

Chapter 6, "Designing a Vulnerability Management Program," and Chapter 7, "Analyzing Vulnerability Scans," of this book focus extensively on the identification and management of vulnerabilities.

Determine Likelihood, Impact, and Risk

After identifying the threats and vulnerabilities facing an organization, risk assessors next seek out combinations of threat and vulnerability that pose a risk to the confidentiality, integrity, or availability of enterprise information and systems. This requires assessing both the likelihood that a risk will materialize and the impact that the risk will have on the organization if it does occur.

When determining the likelihood of a risk occurring, analysts should consider two factors. First, they should assess the likelihood that the threat source will initiate the risk. In the case of an adversarial threat source, this is the likelihood that the adversary will execute an attack against the organization. In the case of accidental, structural, or environmental threats, it is the likelihood that the threat will occur. The second factor that contributes is the likelihood that, if a risk occurs, it will actually have an adverse impact on the organization,

given the state of the organization's security controls. After considering each of these criteria, risk assessors assign an overall likelihood rating. This may use categories, such as "low," "medium," and "high," to describe the likelihood qualitatively.

Risk assessors evaluate the impact of a risk using a similar rating scale. This evaluation should assume that a threat does take place and causes a risk to the organization and then attempt to identify the magnitude of the adverse impact that the risk will have on the organization. When evaluating this risk, it is helpful to refer to the three objectives of cybersecurity shown in Figure 1.1, confidentiality, integrity, and availability, and then assess the impact that the risk would have on each of these objectives.

Exam Note

The risk assessment process described here, using categories of "high," "medium," and "low," is an example of a qualitative risk assessment process. Risk assessments also may use quantitative techniques that numerically assess the likelihood and impact of risks. Quantitative risk assessments are beyond the scope of the Cybersecurity Analyst (CySA+) exam but are found on more advanced security exams, including the CompTIA Advanced Security Practitioner (CASP+) and Certified Information Systems Security Professional (CISSP) exams.

After assessing the likelihood and impact of a risk, risk assessors then combine those two evaluations to determine an overall risk rating. This may be as simple as using a matrix similar to the one shown in Figure 1.4 that describes how the organization assigns overall ratings to risks. For example, an organization might decide that the likelihood of a hacker attack is medium whereas the impact would be high. Looking this combination up in Figure 1.4 reveals that it should be considered a high overall risk. Similarly, if an organization assesses the likelihood of a flood as medium and the impact as low, a flood scenario would have an overall risk of low.

FIGURE 1.4 Many organizations use a risk matrix to determine an overall risk rating based on likelihood and impact assessments.

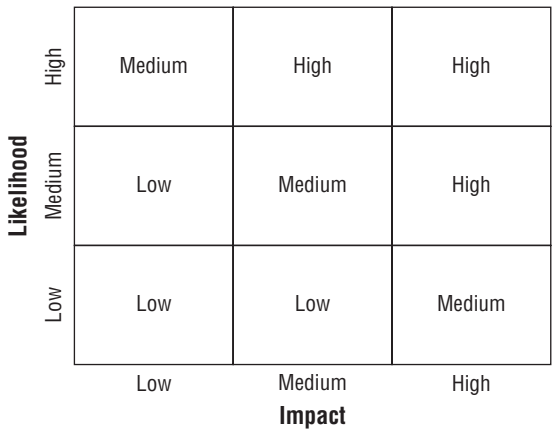

Reviewing Controls

Cybersecurity professionals use risk management strategies, such as risk acceptance, risk avoidance, risk mitigation, and risk transference, to reduce the likelihood and impact of risks identified during risk assessments. The most common way that organizations manage security risks is to develop sets of technical and operational security controls that mitigate those risks to acceptable levels.

Technical controls are systems, devices, software, and settings that work to enforce confidentiality, integrity, and/or availability requirements. Examples of technical controls include building a secure network and implementing endpoint security, two topics discussed later in this chapter. Operational controls are practices and procedures that bolster cybersecurity. Examples of operational controls include conducting penetration testing and using reverse engineering to analyze acquired software. These two topics are also discussed later in this chapter.

Building a Secure Network

Many threats to an organization's cybersecurity exploit vulnerabilities in the organization's network to gain initial access to systems and information. To help mitigate these risks, organizations should focus on building secure networks that keep attackers at bay. Examples of the controls that an organization may use to contribute to building a secure network include network access control (NAC) solutions; network perimeter security controls, such as firewalls; network segmentation; and the use of deception as a defensive measure.

Exam Note

Much of the material in this chapter is not directly testable on the CySA+ exam. You won't, for example, find a question asking you how to configure a NAC solution. However, that doesn't mean that you don't need to know this material! You'll need to be familiar with the security controls in this chapter in order to analyze logs, recommend remediations, and conduct many other activities that *are* directly testable on the exam. The exam does assume that you have a basic familiarity with cybersecurity tools and techniques.

Network Access Control

One of the basic security objectives set forth by most organizations is controlling access to the organization's network. Network access control (NAC) solutions help security professionals achieve two cybersecurity objectives: limiting network access to authorized

individuals and ensuring that systems accessing the organization's network meet basic security requirements.

The 802.1X protocol is a common standard used for NAC. When a new device wishes to gain access to a network, either by connecting to a wireless access point or plugging into a wired network port, the network challenges that device to authenticate using the 802.1X protocol. A special piece of software, known as a supplicant, resides on the device requesting to join the network. The supplicant communicates with a service known as the *authenticator* that runs on either the wireless access point or the network switch. The authenticator does not have the information necessary to validate the user itself, so it passes access requests along to an authentication server using the Remote Authentication Dial-In User Service (RADIUS) protocol. If the user correctly authenticates and is authorized to access the network, the switch or access point then joins the user to the network. If the user does not successfully complete this process, the device is denied access to the network or may be assigned to a special quarantine network for remediation. Figure 1.5 shows the devices involved in 802.1X authentication.

FIGURE 1.5 In an 802.1X system, the device attempting to join the network runs a NAC supplicant, which communicates with an authenticator on the network switch or wireless access point. The authenticator uses RADIUS to communicate with an authentication server.

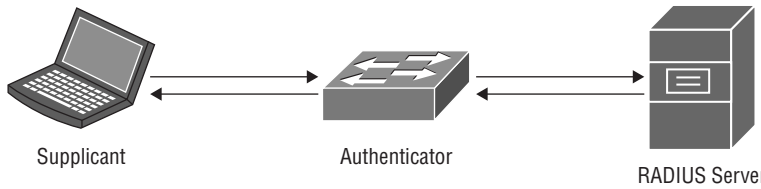

Many NAC solutions are available on the market, and they differ in two major ways:

Agent-Based vs. Agentless Agent-based solutions, such as 802.1X, require that the device requesting access to the network run special software designed to communicate with the NAC service. Agentless approaches to NAC conduct authentication in the web browser and do not require special software.

In-Band vs. Out-of-Band In-band (or inline) NAC solutions use dedicated appliances that sit in between devices and the resources that they wish to access. They deny or limit network access to devices that do not pass the NAC authentication process. The "captive portal" NAC solutions found in hotels that hijack all web requests until the guest enters a room number are examples of in-band NAC. Out-of-band NAC solutions, such as 802.1X, leverage the existing network infrastructure and have network devices communicate with authentication servers and then reconfigure the network to grant or deny network access, as needed.

NAC solutions are often used simply to limit access to authorized users based on those users successfully authenticating, but they may also make network admission decisions based on other criteria. Some of the criteria used by NAC solutions are as follows:

Time of Day Users may be authorized to access the network only during specific time periods, such as during business hours.

Role Users may be assigned to particular network segments based on their role in the organization. For example, a college might assign faculty and staff to an administrative network that may access administrative systems while assigning students to an academic network that does not allow such access.

Location Users may be granted or denied access to network resources based on their physical location. For example, access to the datacenter network may be limited to systems physically present in the datacenter.

System Health NAC solutions may use agents running on devices to obtain configuration information from the device. Devices that fail to meet minimum security standards, such as having incorrectly configured host firewalls, outdated virus definitions, or missing security patches, may be either completely denied network access or placed on a special quarantine network where they are granted only the limited access required to update the system's security.

Administrators may create NAC rules that limit access based on any combination of these characteristics. NAC products provide the flexibility needed to implement the organization's specific security requirements for network admission.

You'll sometimes see the acronym NAC expanded to "Network Admission Control" instead of "network access control." In both cases, people are referring to the same general technology. Network Admission Control is a proprietary name used by Cisco for its network access control solutions.

Firewalls and Network Perimeter Security

NAC solutions are designed to manage the systems that connect directly to an organization's wired or wireless network. They provide excellent protection against intruders who seek to gain access to the organization's information resources by physically accessing a facility and connecting a device to the physical network. They don't provide protection against intruders seeking to gain access over a network connection. That's where firewalls enter the picture.

Network firewalls sit at the boundaries between networks and provide perimeter security. Much like a security guard might control the physical perimeter of a building, the network firewall controls the electronic perimeter. Firewalls are typically configured in the triple-homed fashion illustrated in Figure 1.6. Triple-homed simply means that the firewall connects to three different networks. The firewall in Figure 1.6 connects to the Internet,

the internal network, and a special network known as the demilitarized zone (DMZ), or screened subnet. Any traffic that wishes to pass from one zone to another, such as between the Internet and the internal network, must pass through the firewall.

FIGURE 1.6 A triple-homed firewall connects to three different networks, typically an internal network, a screened subnet, and the Internet.

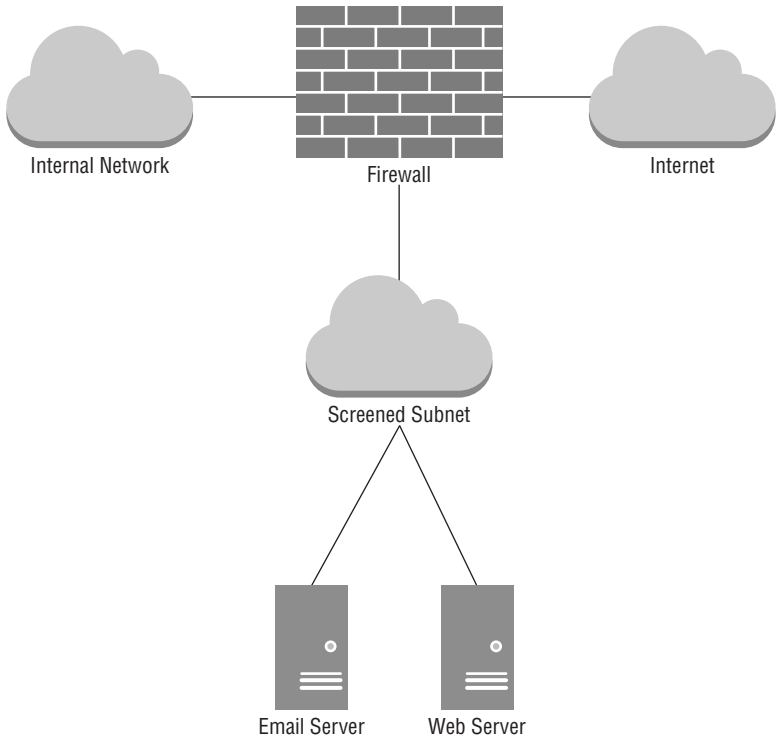

The screened subnet is a special network zone designed to house systems that receive connections from the outside world, such as web and email servers. Sound firewall designs place these systems on an isolated network where, if they become compromised, they pose little threat to the internal network because connections between the screened subnet and the internal network must still pass through the firewall and are subject to its security policy.

Whenever the firewall receives a connection request, it evaluates it according to the firewall's rule base. This rule base is an access control list (ACL) that identifies the types of traffic permitted to pass through the firewall. The rules used by the firewall typically specify the source and destination IP addresses for traffic as well as the destination port corresponding to the authorized service. A list of common ports appears in Table 1.1. Firewalls follow the default deny principle, which says that if there is no rule explicitly allowing a connection, the firewall will deny that connection.

TABLE 1.1 Common TCP ports

Port	Service
20,21	FTP
22	SSH
23	Telnet
25	SMTP
53	DNS
80	HTTP
110	POP3
123	NTP
143	IMAP
389	LDAP
443	HTTPS
636	LDAPS
1433	SQL Server
1521	Oracle
1723	PPTP
3389	RDP

Several categories of firewalls are available on the market today, and they vary in both price and functionality:

- Packet filtering firewalls simply check the characteristics of each packet against the firewall rules without any additional intelligence. Packet filtering firewall capabilities are typically found in routers and other network devices and are very rudimentary firewalls.

- Stateful inspection firewalls go beyond packet filters and maintain information about the state of each connection passing through the firewall. These are the most basic firewalls sold as stand-alone products.
- Next-generation firewalls (NGFWs) incorporate even more information into their decision-making process, including contextual information about users, applications, and business processes. They are the current state-of-the-art in network firewall protection and are quite expensive compared to stateful inspection devices.
- Web application firewalls (WAFs) are specialized firewalls designed to protect against web application attacks, such as SQL injection and cross-site scripting.

Network Segmentation

Firewalls use a principle known as *network segmentation* to separate networks of differing security levels from each other. This principle certainly applies to the example shown in Figure 1.6, where the internal network, screened subnet, and Internet all have differing security levels. The same principle may be applied to further segment the internal network into different zones of trust.

For example, imagine an organization that has several hundred employees and a large datacenter located in its corporate headquarters. The datacenter may house many sensitive systems, such as database servers that contain sensitive employee information, business plans, and other critical information assets. The corporate network may house employees, temporary contractors, visitors, and other people who aren't entirely trusted. In this common example, security professionals would want to segment the datacenter network so that it is not directly accessible by systems on the corporate network. This can be accomplished using a firewall, as shown in Figure 1.7.

The network shown in Figure 1.7 uses a triple-homed firewall, just as was used to control the network perimeter with the Internet in Figure 1.6. The concept is identical, except in this case the firewall is protecting the perimeter of the datacenter from the less trusted corporate network.

Notice that the network in Figure 1.7 also contains a screened subnet with a server called the *jump box*. The purpose of this server is to act as a secure transition point between the corporate network and the datacenter network, providing a trusted path between the two zones. System administrators who need to access the datacenter network should not connect devices directly to the datacenter network but should instead initiate an administrative connection to the jump box, using Secure Shell (SSH), the Remote Desktop Protocol (RDP), or a similar secure remote administration protocol. After successfully authenticating to the jump box, they may then connect from the jump box to the datacenter network, providing some isolation between their own systems and the datacenter network. Connections to the jump box should be carefully controlled and protected with strong multifactor authentication (MFA) technology.

FIGURE 1.7 A triple-homed firewall may also be used to isolate internal network segments of varying trust levels.

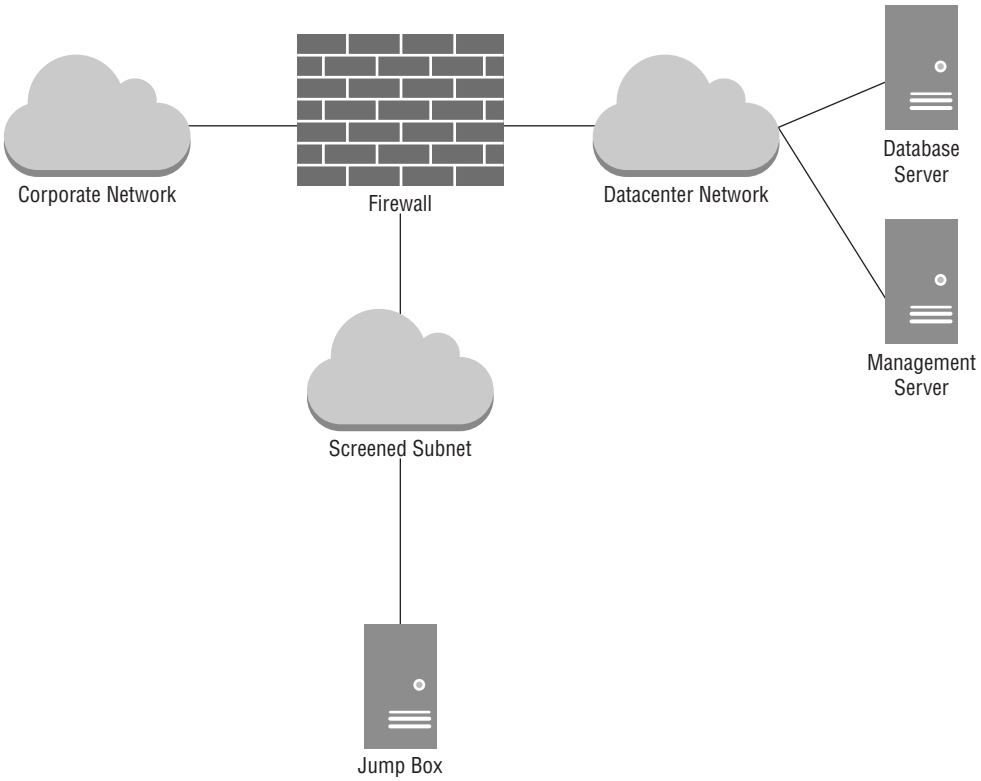

Jump boxes may also be used to serve as a layer of insulation against systems that may only be partially trusted. For example, if you have contractors who bring equipment owned by their employer onto your network or employees bringing personally owned devices, you might use a jump box to prevent those systems from directly connecting to your company's systems.

Defense Through Deception

Cybersecurity professionals may wish to go beyond typical security controls and engage in active defensive measures that actually lure attackers to specific targets and seek to monitor their activity in a carefully controlled environment.

Honeypots are systems designed to appear to attackers as lucrative targets due to the services they run, vulnerabilities they contain, or sensitive information that they appear to host. The reality is that honeypots are designed by cybersecurity experts to falsely appear vulnerable and fool malicious individuals into attempting an attack against them. When an

attacker tries to compromise a honeypot, the honeypot simulates a successful attack and then monitors the attacker's activity to learn more about their intentions. Honeypots may also be used to feed network blacklists, blocking all inbound activity from any IP address that attacks the honeypot.

DNS sinkholes feed false information to malicious software that works its way onto the enterprise network. When a compromised system attempts to obtain information from a DNS server about its command-and-control (C&C or C2) server, the DNS server detects the suspicious request and, instead of responding with the correct answer, responds with the IP address of a sinkhole system designed to detect and remediate the botnet-infected system.

Secure Endpoint Management

Laptop and desktop computers, tablets, smartphones, and other endpoint devices are a constant source of security threats on a network. These systems interact directly with end users and require careful configuration management to ensure that they remain secure and do not serve as the entry point for a security vulnerability on enterprise networks. Fortunately, by taking some simple security precautions, technology professionals can secure these devices against most attacks.

Hardening System Configurations

Operating systems are extremely complex pieces of software designed to perform thousands of different functions. The large code bases that make up modern operating systems are a frequent source of vulnerabilities, as evidenced by the frequent security patches issued by operating system vendors.

One of the most important ways that system administrators can protect endpoints is by hardening their configurations, making them as attack-resistant as possible. This includes disabling any unnecessary services or ports on the endpoints to reduce their susceptibility to attack, ensuring that secure configuration settings exist on devices and centrally controlling device security settings.

Patch Management

System administrators must maintain current security patch levels on all operating systems and applications under their care. Once the vendor releases a security patch, attackers are likely already aware of a vulnerability and may immediately begin preying on susceptible systems. The longer an organization waits to apply security patches, the more likely it becomes that they will fall victim to an attack. That said, enterprises should always test patches prior to deploying them on production systems and networks.

Fortunately, patch management software makes it easy to centrally distribute and monitor the patch level of systems throughout the enterprise. For example, Microsoft's Endpoint

Manager allows administrators to quickly view the patch status of enterprise systems and remediate any systems with missing patches.

Compensating Controls

In some cases, security professionals may not be able to implement all of the desired security controls due to technical, operational, or financial reasons. For example, an organization may not be able to upgrade the operating system on retail point-of-sale (POS) terminals due to an incompatibility with the POS software. In these cases, security professionals should seek out compensating controls designed to provide a similar level of security using alternate means. In the POS example, administrators might place the POS terminals on a segmented, isolated network and use intrusion prevention systems to monitor network traffic for any attempt to exploit an unpatched vulnerability and block it from reaching the vulnerable host. This meets the same objective of protecting the POS terminal from compromise and serves as a compensating control.

Group Policies

Group Policies provide administrators with an efficient way to manage security and other system configuration settings across a large number of devices. Microsoft's Group Policy Object (GPO) mechanism allows administrators to define groups of security settings once and then apply those settings to either all systems in the enterprise or a group of systems based on role.

For example, Figure 1.8 shows a GPO designed to enforce Windows Firewall settings on sensitive workstations. This GPO is configured to require the use of Windows Firewall and block all inbound connections.

Administrators may use GPOs to control a wide variety of Windows settings and create different policies that apply to different classes of systems.

Endpoint Security Software

Endpoint systems should also run specialized security software designed to enforce the organization's security objectives. At a minimum, this should include antivirus software designed to scan the system for signs of malicious software that might jeopardize the security of the endpoint. Administrators may also choose to install host-based firewall software that serves as a basic firewall for that individual system, complementing network-based firewall controls or host intrusion prevention systems (HIPSs) that block suspicious network activity. Endpoint security software should report its status to a centralized management system that allows security administrators to monitor the entire enterprise from a single location.

FIGURE 1.8 Group Policy Objects (GPOs) may be used to apply settings to many different systems at the same time.

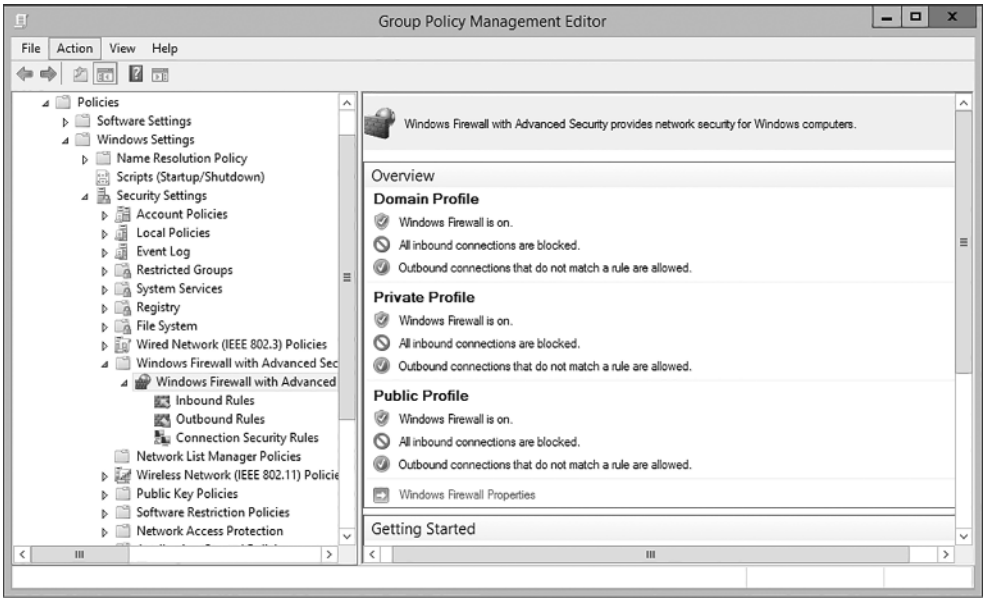

Mandatory Access Controls

In highly secure environments, administrators may opt to implement a mandatory access control (MAC) approach to security. In a MAC system, administrators set all security permissions, and end users cannot modify those permissions. This stands in contrast to the discretionary access control (DAC) model found in most modern operating systems where the owner of a file or resource controls the permissions on that resource and can delegate them at their discretion.

MAC systems are very unwieldy and, therefore, are rarely used outside of very sensitive government and military applications. Security-Enhanced Linux (SELinux), an operating system developed by the U.S. National Security Agency, is an example of a system that enforces mandatory access controls.

Penetration Testing

In addition to bearing responsibility for the design and implementation of security controls, cybersecurity analysts are responsible for monitoring the ongoing effectiveness of those controls. Penetration testing is one of the techniques they use to fulfill this obligation. During

a penetration test, the testers simulate an attack against the organization using the same information, tools, and techniques available to real attackers. They seek to gain access to systems and information and then report their findings to management. The results of penetration tests may be used to bolster an organization's security controls.

Penetration tests may be performed by an organization's internal staff or by external consultants. In the case of internal tests, they require highly skilled individuals and are quite time-consuming. External tests mitigate these concerns but are often quite expensive to conduct. Despite these barriers to penetration tests, organizations should try to perform them periodically since a well-designed and well-executed penetration test is one of the best measures of an organization's cybersecurity posture.

NIST divides penetration testing into the four phases shown in Figure 1.9.

FIGURE 1.9 NIST divides penetration testing into four phases.

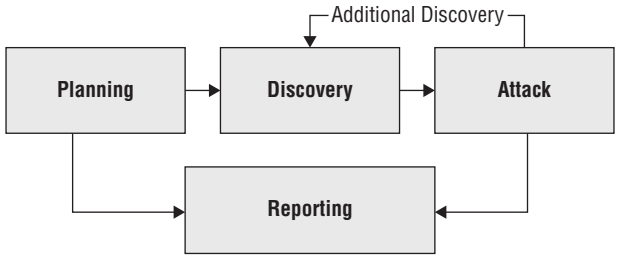

Source: NIST SP 800-115 / U.S Department of Commerce / Public Domain

Planning a Penetration Test

The planning phase of a penetration test lays the administrative groundwork for the test. No technical work is performed during the planning phase, but it is a critical component of any penetration test. There are three important rules of engagement to finalize during the planning phase:

Timing When will the test take place? Will technology staff be informed of the test? Can it be timed to have as little impact on business operations as possible?

Scope What is the agreed-on scope of the penetration test? Are any systems, networks, personnel, or business processes off-limits to the testers?

Authorization Who is authorizing the penetration test to take place? What should testers do if they are confronted by an employee or other individual who notices their suspicious activity?

These details are administrative in nature, but it is important to agree on them up front and in writing to avoid problems during and after the penetration test.

You should never conduct a penetration test without permission. Not only is an unauthorized test unethical, it may be illegal.

Conducting Discovery

The technical work of the penetration test begins during the discovery phase when attackers conduct reconnaissance and gather as much information as possible about the targeted network, systems, users, and applications. This may include conducting reviews of publicly available material, performing port scans of systems, using network vulnerability scanners and web application testers to probe for vulnerabilities, and performing other information gathering.

Vulnerability scanning is an important component of penetration testing. This topic is covered extensively in Chapters 6 and 7.

Executing a Penetration Test

During the attack phase, penetration testers seek to bypass the organization's security controls and gain access to systems and applications run by the organization. Testers often follow the NIST attack process shown in Figure 1.10.

FIGURE 1.10 The attack phase of a penetration test uses a cyclical process that gains a foothold and then uses it to expand access within the target organization.

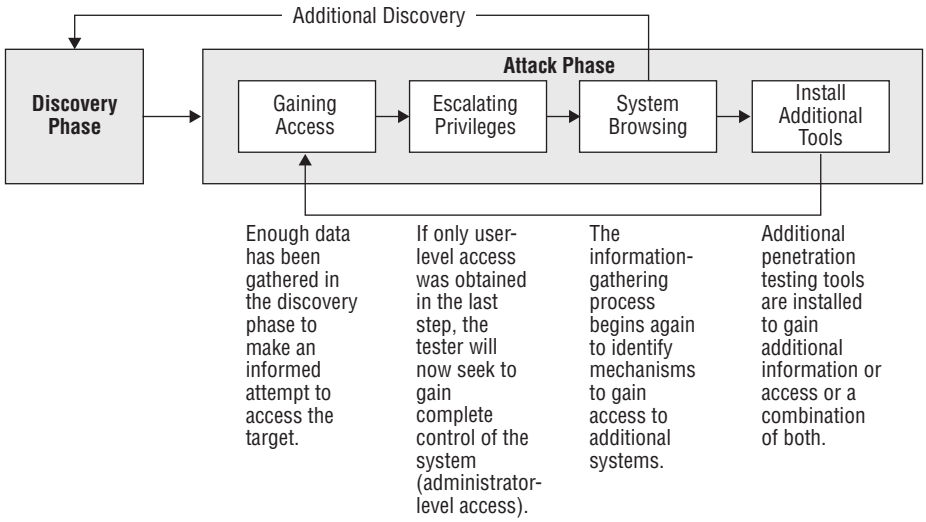

Source: NIST SP 800-115: Technical Guide to Information Security Testing and Assessment

In this process, attackers use the information gathered during the discovery phase to gain initial access to a system. Once they establish a foothold, they then seek to escalate their access until they gain complete administrative control of the system. From there, they can scan for additional systems on the network, install additional penetration testing tools, and begin the cycle anew, seeking to expand their footprint within the targeted organization. They continue this cycle until they exhaust the possibilities or the time allotted for the test expires.

If you're interested in penetration testing, CompTIA offers a companion certification to the CySA+ called PenTest+. The PenTest+ certification dives deeply into penetration testing topics.

Communicating Penetration Test Results

At the conclusion of the penetration test, the testers prepare a detailed report communicating the access they were able to achieve and the vulnerabilities they exploited to gain this access. The results of penetration tests are valuable security planning tools, because they describe the actual vulnerabilities that an attacker might exploit to gain access to a network. Penetration testing reports typically contain detailed appendixes that include the results of various tests and may be shared with system administrators responsible for remediating issues.

Training and Exercises

In addition to performing penetration tests, some organizations choose to run wargame exercises that pit teams of security professionals against one another in a cyberdefense scenario. These exercises are typically performed in simulated environments, rather than on production networks, and seek to improve the skills of security professionals on both sides by exposing them to the tools and techniques used by attackers. Three teams are involved in most cybersecurity wargames:

- The red team plays the role of the attacker and uses reconnaissance and exploitation tools to attempt to gain access to the protected network. The red team's work is similar to that of the testers during a penetration test.
- The blue team is responsible for securing the targeted environment and keeping the red team out by building, maintaining, and monitoring a comprehensive set of security controls.
- The white team coordinates the exercise and serves as referees, arbitrating disputes between the team, maintaining the technical environment, and monitoring the results.

Cybersecurity wargames can be an effective way to educate security professionals on modern attack and defense tactics.

Reverse Engineering

In many cases, vendors do not release the details of how hardware and software work. Certainly, the authors of malicious software don't explain their work to the world. In these situations, security professionals may be in the dark about the security of their environments. Reverse engineering is a technique used to work backward from a finished product to figure out how it works. Security professionals sometimes use reverse engineering to learn the inner workings of suspicious software or inspect the integrity of hardware. Reverse engineering uses a philosophy known as decomposition where reverse engineers start with the finished product and work their way back to its component parts.

Isolation and Sandboxing

One of the most dangerous threats to the security of modern organizations is customized malware developed by APT actors who create specialized tools designed to penetrate a single target. Since they have never been used before, these tools are not detectable with the signature-detection technology used by traditional antivirus software.

Sandboxing is an approach used to detect malicious software based on its behavior rather than its signatures. Sandboxing systems watch systems and the network for unknown pieces of code and, when they detect an application that has not been seen before, immediately isolate that code in a special environment known as a *sandbox* where it does not have access to any other systems or applications. The sandboxing solution then executes the code and watches how it behaves, checking to see if it begins scanning the network for other systems, gathering sensitive information, communicating with a command-and-control server, or performing any other potentially malicious activity.

If the sandboxing solution identifies strange behavior, it blocks the code from entering the organization's network and flags it for administrator review. This process, also known as *code detonation*, is an example of an automated reverse engineering technique that takes action based on the observed behavior of software.

Reverse Engineering Software

In most programming languages, developers write software in a human-readable language such as C/C++, Java, Ruby, or Python. Depending on the programming language, the computer may process this code in one of two ways. In interpreted languages, such as Ruby and Python, the computer works directly from the source code. Reverse engineers seeking to analyze code written in interpreted languages can simply read through the code and often get a good idea of what the code is attempting to accomplish.

In compiled languages, such as Java and C/C++, the developer uses a tool called a *compiler* to convert the source code into binary code that is readable by the computer. This binary code is what is often distributed to users of the software, and it is very difficult, if not

impossible, to examine binary code and determine what it is doing, making the reverse engineering of compiled languages much more difficult. Technologists seeking to reverse-engineer compiled code have two options. First, they can attempt to use a specialized program known as a decompiler to convert the binary code back to source code. Unfortunately, however, this process usually does not work very well. Second, they can use a specialized environment and carefully monitor how software responds to different inputs in an attempt to discover its inner workings. In either case, reverse engineering compiled software is extremely difficult.

Real World Scenario

Fingerprinting Software

Although it is difficult to reverse-engineer compiled code, technologists can easily detect whether two pieces of compiled code are identical or whether one has been modified. *Hashing* is a mathematical technique that analyzes a file and computes a unique fingerprint, known as a message digest or hash, for that file. Analysts using hash functions, such as the Secure Hash Algorithm (SHA), can compute the hashes of two files and compare the output values. If the hashes are identical, the file contents are identical. If the hashes differ, the two files contain at least one difference.

Reverse Engineering Hardware

Reverse engineering hardware is even more difficult than reverse engineering software because the authenticity of hardware often rests in the invisible code embedded within integrated circuits and firmware contents. Although organizations may perform a physical inspection of hardware to detect tampering, it is important to verify that hardware has source authenticity, meaning that it comes from a trusted, reliable source, because it is simply too difficult to exhaustively test hardware.

The U.S. government recognizes the difficulty of ensuring source authenticity and operates a trusted foundry program for critical defense systems. The U.S. Department of Defense (DoD) and National Security Agency (NSA) certify companies as trusted foundries that are approved to create sensitive integrated circuits for government use. Companies seeking trusted foundry status must show that they completely secure the production process, including design, prototyping, packing, assembly, and other elements of the process.

Reverse engineers seeking to determine the function of hardware use some of the same techniques used for compiled software, particularly when it comes to observing behavior. Operating a piece of hardware in a controlled environment and observing how it responds to different inputs provides clues to the functions performed in the hardware. Reverse engineers may also seek to obtain documentation from original equipment manufacturers (OEMs) that provide insight into how components of a piece of hardware function.

Real World Scenario

Compromising Cisco Routers

According to NSA documents released by Edward Snowden, the U.S. government has engaged in reverse engineering of hardware designed to circumvent security.

(TS//SI//NF) Left: Intercepted packages are opened carefully; Right: A "load station" implants a beacon.

Source: Electronic Frontier Foundation / CC BY 3.0 US.

In a process shown in this photo, NSA employees intercepted packages containing Cisco routers, switches, and other network gear after it left the factory and before it reached the customer. They then opened the packages and inserted covert firmware into the devices that facilitated government monitoring.

Efficiency and Process Improvement

Cybersecurity analysts perform a tremendous amount of routine work. From analyzing logs to assessing vulnerabilities, our work can be sometimes tedious and often error-prone. Strong cybersecurity teams invest time and money in improving their own processes, using standardization and automation to improve their efficiency, reduce the likelihood of error, and free up the valuable time of analysts for more significant work.

Exam Note

As you prepare for the exam, you should be very focused on this content, as it is directly listed in the exam objectives. CompTIA wants cybersecurity analysts to focus on standardizing processes and streamlining operations. A very effective way to achieve this goal is to integrate diverse technologies and tools into a "single pane of glass" monitoring solution.

Standardize Processes and Streamline Operations

One of the first ways that teams can improve their efficiency is to create standardized processes for recurring activities. When you find yourself facing a task where you're inventing a solution by the seat of your pants, that's a good sign that you might benefit from taking the time to create a standardized process for handling similar events in the future. This is especially true for activities you find yourself repeating time and time again.

Standardizing processes reduces the amount of effort required to react to a task. You no longer need to figure out what to do the next time this task comes up—you simply turn to your playbook for that standardized process and carry out the steps that you've already thought through. These processes also have the added benefit of ensuring that different members of the team respond consistently in similar situations.

Cybersecurity Automation

Standardizing tasks also helps you identify opportunities for automation. You may be able to go beyond standardizing the work of team members and automate some responses to take people out of the loop entirely. In the terms used by CompTIA, you're "minimizing human engagement." *Security orchestration, automation, and response (SOAR)* platforms provide many opportunities to automate security tasks that cross between multiple systems. You may wish to coordinate with other members of your team taking an inventory of all the activities performed by the team and identify those that are suitable for automation. The two key characteristics of processes that can be automated are that they are both repeatable and do not require human interaction. Once you have automations in place, you'll just need to coordinate with your team to manage existing automations and facilitate the adoption of new automations.

SOAR platforms also offer opportunities to improve your organization's use of threat intelligence. By bringing information about emerging threats into your SOAR platform, you can enrich data about ongoing incidents and improve your ability to react to emerging cybersecurity situations. The SOAR platform provides you with the opportunity to combine information received through multiple threat feeds and develop a comprehensive picture of the cybersecurity landscape and your security posture.

Technology and Tool Integration

Cybersecurity is full of tedious work. From performing vulnerability scans of large networks to conducting file integrity tests, we often need to use very repetitive processes to achieve our objectives. If we tried to get this work done manually, it would be so time-consuming that we would stop from exhaustion before we ever finished.

Workflow orchestration techniques help us automate many of these repetitive tasks. In addition to the SOAR platforms we've already discussed, there are two more ways that we can automate our workflows: *scripting*, which is writing code that automates our work, and *integration*, which uses vendor-provided interfaces to tie different products together.

Analysts commonly take advantage of *application programming interfaces (APIs)* as a primary means of integrating diverse security tools. APIs are programmatic interfaces to services that allow you to interact with that service without using web-based interfaces. You can normally perform the same actions with an API that you could perform at a service's web-based interface, but the API allows you to write code to automate those actions.

APIs are powerful to cybersecurity analysts because we can use them to reach into a wide variety of systems. Our security tools often offer APIs that we can leverage in our automation work, but so do many other technology services. We can use APIs to automate the provisioning of cloud resources, to retrieve access logs from remote services, and automate many other routine tasks.

Webhooks allow us to send a signal from one application to another using a web request. For example, you might want to run a vulnerability scan every time your threat intelligence platform receives a report of a new vulnerability. In that case, you may be able to configure a webhook action in the threat intelligence platform that sends a request to the vulnerability scanner's API each time a new vulnerability is reported. That request could trigger the desired scan.

Plug-ins also provide an opportunity to increase integrations. These are small programs that run inside of other programs, adding additional functionality. For example, you might use a browser plug-in to perform data enrichment. That plug-in might automatically pull up Whois and reputation data each time you hover over an IP address in your browser.

The ultimate goal of cybersecurity automation is to achieve a *single pane of glass* approach to security operations. In this philosophy, cybersecurity analysts integrate all their tools into a single platform, so they can use one consistent interface to perform all of their work. Now, of course, this isn't really an achievable goal. There's always going to be "one more system" that prevents you from truly reaching a single pane of glass approach. However, as a design principle, reducing the number of interfaces that cybersecurity analysts must use each day can dramatically increase their efficiency.

Bringing Efficiency to Incident Response

Incident response is one of the rapidly emerging areas of automation, as security teams seek to bring the power of automation to what is often the most human-centric task in cybersecurity: investigating anomalous activity. While SOAR automation and other security tools may

trigger an incident investigation, the work of the incident responder from that point forward is often a very manual process that involves the application of tribal knowledge, personal experience, and instinct.

Enriching Incident Response Data

While incident response will likely always involve a significant component of human intervention, some organizations are experiencing success with automating portions of their incident response programs. One of the best starting points for incident response automation involves providing automated enrichment of incident response data to human analysts, saving them the tedious time of investigating routine details of an incident. For example, when an intrusion detection system identifies a potential attack, a security automation workflow can trigger a series of activities. These might include:

- Performing reconnaissance on the source address of the attack, including IP address ownership and geolocation information
- Supplementing the initial report with other log information for the targeted system based upon a security information and event management (SIEM) query. SIEMs are an important security technology that you'll find covered in Chapter 3, "Malicious Activity."
- Triggering a vulnerability scan of the targeted system designed to assist in determining whether the attack has a high likelihood of success

All of these actions can take place immediately upon the detection of the incident and appended to the incident report in the tracking system for review by a cybersecurity analyst. Teams seeking to implement incident response data enrichment will benefit from observing the routine activities of first responders and identifying any information gathering requirements that are possible candidates for automation.

Automate Portions of Incident Response Playbooks

In addition to enriching the data provided to cybersecurity analysts responding to an incident, automation may also play a role in improving the efficiency of response actions. Incident response teams commonly use playbooks that define the sequence of steps they will follow when responding to common incident types. Actions defined in a playbook may include activities designed to contain the incident, eradicate the effects of the incident from the network, and recover normal operations. You'll find more coverage of playbooks in Chapter 9, "Building an Incident Response Program."

In some cases, it is reasonable to believe that the response to a security incident may be fully automated. For example, if a system on the local network begins emitting high volumes of traffic targeted at remote web servers, a reasonable analyst might conclude that the system is compromised and participating in a botnet-driven denial-of-service attack. The natural reaction of that analyst might then be to immediately quarantine the system to contain the damage and then send a case to the appropriate IT support technician to investigate and resolve the issue. That straightforward response is a prime candidate for automation.

Integrations between the security automation platform, networking devices, and the ticketing systems can perform all these activities without any human intervention.

While the response to some simple incidents may be fully automated, many incidents have unique characteristics that require some degree of human intervention. Even in those cases, portions of the incident response playbook may still be automated beyond the data enrichment stage. For example, a cybersecurity analyst may have a dashboard that allows them to trigger a sequence of activities that blocks access to all resources for an individual user, quarantines suspect systems, or activates an incident escalation process.

The Future of Cybersecurity Analytics

As we continue to develop our cybersecurity analytics capabilities, the tools and techniques available to us advance in sophistication. The area of greatest promise for future cybersecurity analytics tools is the continued adoption of *machine learning* techniques designed to automatically extract knowledge from the voluminous quantity of information generated by security systems.

Machine learning techniques are already incorporated into many security analytics tools, providing automated analysis of data based on the experiences of other users of the tool. Expect to see these capabilities continue to develop as organizations harvest the power of machine learning to reduce the requirements for human analysts to perform burdensome sifting of data and allow them to focus on the output of machine learning algorithms that guide them toward more productive work.

Summary

Cybersecurity professionals are responsible for ensuring the confidentiality, integrity, and availability of information and systems maintained by their organizations. Confidentiality ensures that unauthorized individuals are not able to gain access to sensitive information. Integrity ensures that there are no unauthorized modifications to information or systems, either intentionally or unintentionally. Availability ensures that information and systems are ready to meet the needs of legitimate users at the time those users request them. Together, these three goals are known as the CIA Triad.

As cybersecurity analysts seek to protect their organizations, they must evaluate risks to the CIA Triad. This includes identifying vulnerabilities, recognizing corresponding threats, and determining the level of risk that results from vulnerability and threat combinations. Analysts must then evaluate each risk and identify appropriate risk management strategies to mitigate or otherwise address the risk. They may use machine learning techniques to assist with this work.

Cybersecurity analysts mitigate risks using security controls designed to reduce the likelihood or impact of a risk. Network security controls include network access control (NAC) systems, firewalls, and network segmentation. Secure endpoint controls include hardened system configurations, patch management, Group Policies, and endpoint security software.

Penetration tests and reverse engineering provide analysts with the reassurance that the controls they've implemented to mitigate risks are functioning properly. By following a careful risk analysis and control process, analysts significantly enhance the confidentiality, integrity, and availability of information and systems under their control.

Finally, cybersecurity analysts should take the time to invest in their own operations practices through efficiency and process improvement initiatives. Projects that standardize processes, streamline operations, and integrate technologies and tools make the work of cybersecurity analysts easier, increasing productivity and reducing the likelihood of error.

Exam Essentials

Know the three objectives of cybersecurity. *Confidentiality* ensures that unauthorized individuals are not able to gain access to sensitive information. *Integrity* ensures that there are no unauthorized modifications to information or systems, either intentionally or unintentionally. *Availability* ensures that information and systems are ready to meet the needs of legitimate users at the time those users request them.

Know how cybersecurity risks result from the combination of a threat and a vulnerability. A vulnerability is a weakness in a device, system, application, or process that might allow an attack to take place. A threat in the world of cybersecurity is an outside force that may exploit a vulnerability.

Be able to categorize cybersecurity threats as adversarial, accidental, structural, or environmental. Adversarial threats are individuals, groups, and organizations that are attempting to deliberately undermine the security of an organization. Accidental threats occur when individuals doing their routine work mistakenly perform an action that undermines security. Structural threats occur when equipment, software, or environmental controls fail due to the exhaustion of resources, exceeding their operational capability or simply failing due to age. Environmental threats occur when natural or human-made disasters occur that are outside the control of the organization.

Understand how networks are made more secure through the use of network access control, firewalls, and segmentation. Network access control (NAC) solutions help security professionals achieve two cybersecurity objectives: limiting network access to authorized individuals and ensuring that systems accessing the organization's network meet basic security requirements. Network firewalls sit at the boundaries between networks and provide perimeter security. Network segmentation uses isolation to separate networks of differing security levels from each other.

Understand how endpoints are made more secure through the use of hardened configurations, patch management, Group Policy, and endpoint security software. Hardening configurations includes disabling any unnecessary services on the endpoints to reduce their susceptibility to attack, ensuring that secure configuration settings exist on devices, and centrally controlling device security settings. Patch management ensures that operating systems and applications are not susceptible to known vulnerabilities. Group Policy allows the application of security settings to many devices simultaneously, and endpoint security software protects against malicious software and other threats.

Know that penetration tests provide organizations with an attacker's perspective on their security. The NIST process for penetration testing divides tests into four phases: planning, discovery, attack, and reporting. The results of penetration tests are valuable security planning tools, since they describe the actual vulnerabilities that an attacker might exploit to gain access to a network.

Understand how reverse engineering techniques attempt to determine how hardware and software function internally. Sandboxing is an approach used to detect malicious software based on its behavior rather than its signatures. Other reverse engineering techniques are difficult to perform, are often unsuccessful, and are quite time-consuming.

Know how machine learning technology facilitates cybersecurity analysis. The area of greatest promise for future cybersecurity analytics tools is the continued adoption of *machine learning* techniques designed to automatically extract knowledge from the voluminous quantity of information generated by security systems. Machine learning techniques are already incorporated into many security analytics tools, providing automated analysis of data based on the experiences of other users of the tool.

Understand the importance of efficiency and process improvements to security operations. Cybersecurity analysis is difficult, and sometimes tedious, work. Projects that work to streamline operations, standardize processes, and increase integrations of tools and technologies allow work to flow more smoothly. This not only boosts the efficiency and effectiveness of cybersecurity analysts but also reduces the likelihood of error.

Lab Exercises

Activity 1.1: Create an Inbound Firewall Rule

In this lab, you will verify that the Windows Defender Firewall is enabled on a server and then create an inbound firewall rule that blocks file and printer sharing.

These lab instructions were written to run on a system running Windows Server 2022. The process for working on other versions of Windows Server is quite similar, although the exact names of services, options, and icons may differ slightly.

You should perform this lab on a test system. Enabling file and printer sharing on a production system may have undesired consequences. The easiest way to get access to a Windows Server 2022 system is to create an inexpensive cloud instance through Amazon Web Services (AWS) or Microsoft Azure.

Part 1: Verify that Windows Defender Firewall is enabled

1. Open Control Panel for your Windows Server.
2. Choose System and Security.
3. Under Windows Defender Firewall, click Check Firewall Status.
4. Verify that the Windows Defender Firewall state is set to On for Private networks. If it is not on, enable the firewall by using the "Turn Windows Defender Firewall on or off" link on the left side of the window.

Part 2: Create an inbound firewall rule that allows file and printer sharing

1. On the left side of the Windows Defender Firewall Control Panel, click "Allow an app or feature through Windows Defender Firewall."
2. Scroll down the list of applications and find File and Printer Sharing.
3. Check the box to the left of that entry to block connections related to File and Printer Sharing.
4. Confirm that the Private box to the right of that option was automatically selected. This allows File and Printer Sharing only for other systems on the same local network. The box for public access should be unchecked, specifying that remote systems are not able to access this feature.
5. Click OK to apply the setting.

Activity 1.2: Create a Group Policy Object

In this lab, you will create a Group Policy Object and edit its contents to enforce an organization's password policy.

These lab instructions were written to run on a system running Windows Server 2022. The process for working on other versions of Windows Server is quite similar, although the exact names of services, options, and icons may differ slightly. To complete this lab, your Windows Server must be configured as a domain controller.

1. Open the Group Policy Management application. (If you do not find this application on your Windows Server, it is likely that it is not configured as a domain controller.)
2. Expand the folder corresponding to your Active Directory forest.
3. Expand the `Domains` folder.
4. Expand the folder corresponding to your domain.
5. Right-click the `Group Policy Objects` folder and select New from the pop-up menu.

6. Name your new GPO **Password Policy** and click OK.
7. Click the `Group Policy Objects` folder.
8. Right-click the new Password Policy GPO and select Edit from the pop-up menu.
9. When Group Policy Management Editor opens, expand the `Policies` folder under the Computer Configuration section.
10. Expand the `Windows Settings` folder.
11. Expand the `Security Settings` folder.
12. Expand the `Account Policies` folder.
13. Click Password Policy.
14. Double-click Maximum Password Age.
15. In the pop-up window, select the Define This Policy Setting check box and set the expiration value to 90 days.
16. Click OK to close the window.
17. Click OK to accept the suggested change to the minimum password age.
18. Double-click the Minimum Password Length option.
19. As in the prior step, click the box to define the policy setting and set the minimum password length to 12 characters.
20. Click OK to close the window.
21. Double-click the Password Must Meet Complexity Requirements option.
22. Click the box to define the policy setting and change the value to Enabled.
23. Click OK to close the window.
24. Click the X to exit Group Policy Management Editor.

You have now successfully created a Group Policy Object that enforces the organization's password policy. You can apply this GPO to users and/or groups as needed.

Activity 1.3: Write a Penetration Testing Plan

For this activity, you will design a penetration testing plan for a test against an organization of your choosing. If you are employed, you may choose to use your employer's network. If you are a student, you may choose to create a plan for a penetration test of your school. Otherwise, you may choose any organization, real or fictitious, of your choice.

Your penetration testing plan should cover the three main criteria required before initiating any penetration test:

- Timing
- Scope
- Authorization

One word of warning: You should not conduct a penetration test without permission of the network owner. This assignment only asks you to design the test on paper.

Activity 1.4: Recognize Security Tools

Match each of the security tools listed in this table with the correct description.

Firewall	Determines which clients may access a wired or wireless network
Decompiler	Creates a unique fingerprint of a file
Antivirus	Filters network connections based on source, destination, and port
NAC	System intentionally created to appear vulnerable
GPO	Attempts to recover source code from binary code
Hash	Scans a system for malicious software
Honeypot	Protects against SQL injection attacks
WAF	Deploys configuration settings to multiple Windows systems