Chapter 1

Understanding Security Fundamentals

CISCO CCNA SECURITY EXAM OBJECTIVES COVERED IN THIS CHAPTER:

✓ **1.1 Common security principles**

- Describe confidentiality, integrity, availability (CIA)
- Identify common security terms
- Identify common network security zones

✓ **1.4 Describe network topologies**

- Campus area network (CAN)
- Cloud, wide area network (WAN)
- Data center
- Small office/home office (SOHO)
- Network security for a virtual environment

Securing a network is no easy task. Daily you probably hear about data disclosures and new network attacks. However, you are not defenseless. By properly implementing the security features available in Cisco routers, switches, and firewalls, you can reduce the risk of a security breach to a manageable level. This book is designed to help you understand the issues, identify your security options, and deploy those options in the correct manner. In the process, the book will prepare you for the Cisco CCNA Security certification, which validates the skills and knowledge required to secure a network using Cisco products.

In this chapter, you will learn the following:

- Common security principles
- Network topologies

Goals of Security

When you're securing a network, several important security principles should guide your efforts. Every security measure you implement should contribute to the achievement of one of three goals. The three fundamentals of security are confidentiality, integrity, and availability (CIA), often referred to as the *CIA triad*.

Most security issues result in a violation of at least one facet of the CIA triad. Understanding these three security principles will help ensure that the security controls and mechanisms implemented protect at least one of these principles.

Every security control that is put into place by an organization fulfills at least one of the security principles of the CIA triad. Understanding how to circumvent these security principles is just as important as understanding how to provide them.

Confidentiality

To ensure *confidentiality*, you must prevent the disclosure of data or information to unauthorized entities. As part of confidentiality, the sensitivity level of data must be determined before putting any access controls in place. Data with a higher sensitivity level will have more access controls in place than data at a lower sensitivity level. Identification, authentication, and authorization can be used to maintain data confidentiality. Encryption is another popular example of a control that provides confidentiality.

Integrity

Integrity, the second part of the CIA triad, ensures that data is protected from unauthorized modification or data corruption. The goal of integrity is to preserve the consistency of data, including data stored in files, databases, systems, and networks.

An access control list (ACL) is an example of a control that helps to provide integrity. Another example is the generation of hash values that can be used to validate data integrity.

Availability

Availability means ensuring that data is accessible when and where it is needed. Only individuals who need access to data should be allowed access to that data. The two main areas where availability is affected are

- When attacks are carried out that disable or cripple a system.
- When service loss occurs during and after disasters. Each system should be assessed on its criticality to organizational operations. Controls are implemented based on each system's criticality level.

Fault-tolerant technologies, such as RAID or redundant sites, are examples of controls that help to improve availability.

Guiding Principles

When managing network security and access to resources, there are some proven principles that should guide your efforts. These concepts have stood the test of time because they contribute to supporting the CIA triad.

Least Privilege/Need-to-Know

The principle of *least privilege* requires that a user or process is given only the minimum access privilege needed to perform a particular task. Its main purpose is to ensure that users only have access to the resources they need and are authorized to perform only the tasks they need to perform. To properly implement the least privilege principle, organizations must identify all users' jobs and restrict users only to the identified privileges.

The *need-to-know* principle is closely associated with the concept of least privilege. Although least privilege seeks to reduce access to a minimum, the need-to-know principle actually defines what the minimums for each job or business function are. Excessive privileges become a problem when a user has more rights, privileges, and permissions than he needs to do his job. Excessive privileges are hard to control in large environments.

A common implementation of the least privilege and need-to-know principles is when a systems administrator is issued both an administrative-level account and a normal user account. In most day-to-day functions, the administrator should use his normal user account. When the systems administrator needs to perform administrative-level tasks, he should use the administrative-level account. If the administrator uses his administrative-level account

while performing routine tasks, he risks compromising the security of the system and user accountability.

Organizational rules that support the principle of least privilege include the following:

- Keep the number of administrative accounts to a minimum.
- Administrators should use normal user accounts when performing routine operations.
- Permissions on tools that are likely to be used by attackers should be as restrictive as possible.

To more easily support the least privilege and need-to-know principles, users should be divided into groups to facilitate the confinement of information to a single group or area. This process is referred to as *compartmentalization*.

Default to No Access

During the authorization process, you should configure an organization's access control mechanisms so that the default level of security is to default to no access. This means that if nothing has been specifically allowed for a user or group, then the user or group will not be able to access the resource. The best security approach is to start with no access and add rights based on a user's need to know and least privilege needed to accomplish daily tasks.

Defense in Depth

A *defense-in-depth* strategy refers to the practice of using multiple layers of security between data and the resources on which it resides and possible attackers. The first layer of a good defense-in-depth strategy is appropriate access control strategies. Access controls exist in all areas of an information systems (IS) infrastructure (more commonly referred to as an *IT infrastructure*), but a defense-in-depth strategy goes beyond access control. It also considers software development security, cryptography, and physical security. Figure 1.1 shows an example of the defense-in-depth concept.

FIGURE 1.1 Defense in depth

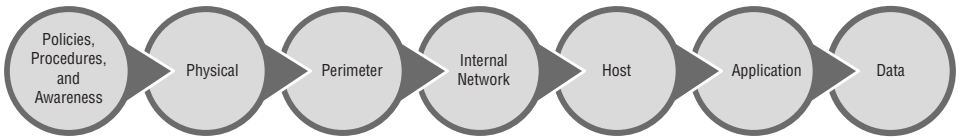

Separation of Duties

Separation of duties is a preventive administrative control to keep in mind when designing an organization's authentication and authorization policies. Separation of duties prevents fraud by distributing tasks and their associated rights and privileges between more than one user. It helps to deter fraud and collusion because when an organization implements adequate separation of duties, collusion between two or more personnel would be required to carry out fraud against the organization. A good example of separation duties is authorizing one person to manage backup procedures and another to manage restore procedures.

Separation of duties is associated with dual controls and split knowledge. With dual controls, two or more users are authorized and required to perform certain functions. For example, a retail establishment might require two managers to open the safe. Split knowledge ensures that no single user has all the information to perform a particular task. An example of a split control is the military requiring two individuals to each enter a unique combination to authorize missile firing.

Separation of duties ensures that one person is not capable of compromising organizational security. Any activities that are identified as high risk should be divided into individual tasks, which can then be allocated to different personnel or departments.

Let's look at an example of the violation of separation of duties. An organization's internal audit department investigates a possible breach of security. One of the auditors interviews three employees.

- A clerk who works in the accounts receivable office and is in charge of entering data into the finance system
- An administrative assistant who works in the accounts payable office and is in charge of approving purchase orders
- The finance department manager who can perform the functions of both the clerk and the administrative assistant

To avoid future security breaches, the auditor should suggest that the manager should only be able to review the data and approve purchase orders.

Job Rotation

From a security perspective, *job rotation* refers to the detective administrative control where multiple users are trained to perform the duties of a position to help prevent fraud by any individual employee. The idea is that by making multiple people familiar with the legitimate functions of the position, the likelihood increases that unusual activities by any one person will be noticed. Job rotation is often used in conjunction with mandatory vacations. Beyond the security aspects of job rotation, additional benefits include the following:

- Trained backup in case of emergencies
- Protection against fraud
- Cross-training of employees

Mandatory Vacation

With *mandatory vacation*s, all personnel are required to take time off, allowing other personnel to fill their positions while gone. This detective administrative control enhances the opportunity to discover unusual activity.

Some of the security benefits of using mandatory vacations include having the replacement employee do the following:

- Run the same applications as the vacationing employee
- Perform tasks in a different order from the vacationing employee
- Perform the job from a different workstation than the vacationing employee

Replacement employees should avoid running scripts that were created by the vacationing employee. A replacement employee should either develop their own script or manually complete the tasks in the script.

Common Security Terms

The risk management process cannot be discussed without understanding some key terms used in risk management. Security professionals should become familiar with the following terms as they are used in risk management:

- *Assets* include anything that is of value to the organization. Assets can be physical such as buildings, land, and computers, and they can be intangible such as data, plans, and recipes.
- A *vulnerability* is an absence or weakness of a countermeasure that is in place. Vulnerabilities can occur in software, hardware, or personnel. An example of a vulnerability is unrestricted access to a folder on a computer. Most organizations implement a vulnerability assessment to identify vulnerabilities.
- A *threat* is the next logical progression in risk management. A threat occurs when vulnerability is identified or exploited. A threat would occur when an attacker identified the folder on the computer that has an inappropriate or absent ACL.
- A *threat agent* is something that carries out a threat. Continuing with the example, the attacker who takes advantage of the inappropriate or absent ACL is the threat agent. Keep in mind, though, that threat agents can discover and/or exploit vulnerabilities. Not all threat agents will actually exploit an identified vulnerability.
- A *risk* is the probability that a threat agent will exploit a vulnerability and the impact if the threat is carried out. The risk in the vulnerability example would be fairly high if the data residing in the folder is confidential. However, if the folder contains only public data, then the risk would be low. Identifying the potential impact of a risk often requires security professionals to enlist the help of subject-matter experts.
- An *exposure* occurs when an organizational asset is exposed to losses. If the folder with the inappropriate or absent ACL is compromised by a threat agent, the organization is exposed to the possibility of data exposure and loss.
- A *countermeasure* reduces the potential risk. Countermeasures are also referred to as *safeguards* or *controls*. Three things must be considered when implementing a countermeasure: vulnerability, threat, and risk. For this example, a good countermeasure would be to implement the appropriate ACL and to encrypt the data. The ACL protects the integrity of the data, and the encryption protects the confidentiality of the data.

Countermeasures or controls come in many categories and types. The categories and types of controls are discussed later in this chapter.

All the aforementioned security concepts work together in the relationship demonstrated in Figure 1.2.

FIGURE 1.2 Security cycle

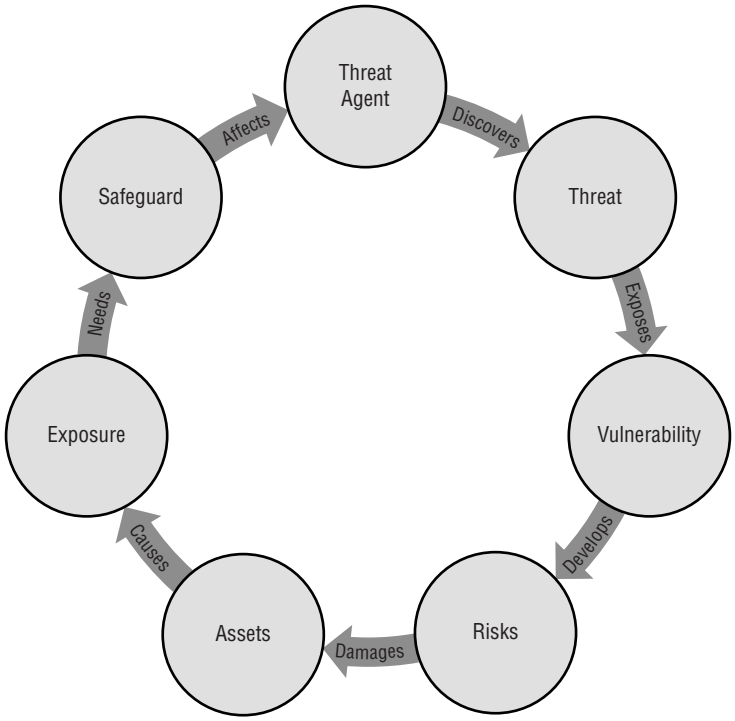

Risk Management Process

The *risk management* process is composed of a series of operations in which the data from one operation feeds the next operation. According to NIST SP 800-30, common information-gathering techniques used in risk analysis include automated risk assessment tools, questionnaires, interviews, and policy document reviews. Keep in mind that multiple sources should be used to determine the risks to a single asset. NIST SP 800-30 identifies the following steps in the risk management process:

1. Identify the assets and their value.
2. Identify threats.
3. Identify vulnerabilities.
4. Determine likelihood.
5. Identify impact.
6. Determine risk as a combination of likelihood and impact.

The following sections include these processes and two additional ones that relate to the identification of countermeasures and cost-benefit analysis.

Asset Classification

The first step of any risk assessment is to identify the assets and determine the asset value, called *asset classification*. Assets are both tangible and intangible. Tangible assets include computers, facilities, supplies, and personnel. Intangible assets include intellectual property, data, and organizational reputation. The value of an asset should be considered in respect to the asset owner's view. The six following considerations can be used to determine the asset's value:

- Value to owner
- Work required developing or obtaining the asset
- Costs to maintain the asset
- Damage that would result if the asset were lost
- Cost that competitors would pay for the asset
- Penalties that would result if the asset was lost

After determining the value of the assets, you should determine the vulnerabilities and threats to each asset.

Data Assets

Data should be classified based on its value to the organization and its sensitivity to disclosure. Assigning a value to data allows an organization to determine the resources that should be used to protect the data. Resources that are used to protect data include personnel resources, monetary resources, access control resources, and so on. Classifying data allows you to apply different protective measures. Data classification is critical to all systems to protect the confidentiality, integrity, and availability of data.

After data is classified, the data can be segmented based on its level of protection needed. The classification levels ensure that data is handled and protected in the most cost-effective manner possible. An organization should determine the classification levels it uses based on the needs of the organization. Several commercial business and military and government information classifications are commonly used.

The information life cycle should also be based on the classification of the data. Organizations are required to retain certain information, particularly financial data, based on local, state, or government laws and regulations.

In this section, we will discuss the sensitivity and criticality of data, commercial business classifications, military and government classifications, information life cycle, database maintenance, and data audit.

SENSITIVITY AND CRITICALITY

Sensitivity is a measure of how freely the data can be handled. Some data requires special care and handling, especially when inappropriate handling could result in penalties, identity theft, financial loss, invasion of privacy, or unauthorized access by an individual or many individuals. Some data is also subject to regulation by state or federal laws and requires notification in the event of a disclosure.

Data is assigned a level of sensitivity based on who should have access to it and how much harm would be done if it were disclosed. This assignment of sensitivity is called *data classification*.

Criticality is a measure of the importance of the data. Data considered sensitive may not necessarily be considered critical. Assigning a level of criticality to a particular data set must take into consideration the answers to a few questions:

- Will you be able to recover the data in case of disaster?
- How long will it take to recover the data?
- What is the effect of this downtime, including loss of public standing?

Data is considered essential when it is critical to the organization's business. When essential data is not available, even for a brief period of time, or its integrity is questionable, the organization will be unable to function. Data is considered required when it is important to the organization, but organizational operations would continue for a predetermined period of time even if the data is not available. Data is nonessential if the organization is able to operate without it during extended periods of time.

Once the sensitivity and criticality of data is understood and documented, the organization should then work to create a data classification system. Most organizations will use either a commercial business classification system or a military and government classification system.

COMMERCIAL BUSINESS CLASSIFICATIONS

Commercial businesses usually classify data using four main classification levels, listed from highest sensitivity level to lowest:

1. Confidential
2. Private
3. Sensitive
4. Public

Data that is confidential includes trade secrets, intellectual data, application programming code, and other data that could seriously affect the organization if unauthorized disclosure occurred. Data at this level would be available only to personnel in the organization whose work relates to the data's subject. Access to confidential data usually requires authorization for each access. Confidential data is exempt from disclosure under the Freedom of Information Act. In most cases, the only way for external entities to have authorized access to confidential data is as follows:

- After signing a confidentiality agreement
- When complying with a court order
- As part of a government project or contract procurement agreement

Data that is private includes any information related to personnel, including human resource records, medical records, and salary information, that is used only within the organization. Data that is sensitive includes organizational financial information and

requires extra measures to ensure its CIA and accuracy. Public data is data that would not cause a negative impact on the organization.

MILITARY AND GOVERNMENT CLASSIFICATIONS

Military and governmental entities usually classify data using five main classification levels, listed from highest sensitivity level to lowest:

1. Top secret
2. Secret
3. Confidential
4. Sensitive but unclassified
5. Unclassified

Data that is top secret includes weapon blueprints, technology specifications, spy satellite information, and other military information that could gravely damage national security if disclosed. Data that is secret includes deployment plans, missile placement, and other information that could seriously damage national security if disclosed. Data that is confidential includes patents, trade secrets, and other information that could seriously affect the government if unauthorized disclosure occurred. Data that is sensitive but unclassified includes medical or other personal data that might not cause serious damage to national security but could cause citizens to question the reputation of the government. Military and government information that does not fall into any of the other four categories is considered unclassified and usually has to be granted to the public based on the Freedom of Information Act.

OTHER CLASSIFICATION SYSTEMS

Another classification system created by the United Kingdom's National Infrastructure Security Coordination Centre (NISCC, now Centre for Protection of National Infrastructure) and since adopted by the ISO/IEC as part of the Standard on Information security management for intersector and interorganizational communications and by CERT is the Traffic Light Protocol (TLP). This system uses traffic light colors to classify information assets. Table 1.1 shows the four colors and their meanings.

TABLE 1.1 TLP classifications

Color	Meaning
Red	Shared only within a meeting
Amber	Shared only with those in the organization with a need to know
Green	Shared only within a community
White	No restriction but still subject to copyright rules

Vulnerability Identification

When identifying vulnerabilities, the Common Vulnerability Scoring System and the Security Content Automation Protocol are standards used in this process. In this section, you'll learn about these two methods for enumerating vulnerabilities in a common format.

Security Content Automation Protocol (SCAP) is a standard used by the security automation community used to enumerate software flaws and configuration issues. It standardized the nomenclature and formats used. A vendor of security automation products can obtain a validation against SCAP, demonstrating that it will interoperate with other scanners and express the scan results in a standardized way.

Understanding the operation of SCAP requires an understanding of the components of it.

Common Configuration Enumeration (CCE) These are configuration best-practice statements maintained by NIST.

Common Platform Enumeration (CPE) These are methods for describing and classifying operating systems applications and hardware devices.

Common Weakness Enumeration (CWE) These are design flaws in the development of software that can lead to vulnerabilities.

Common Vulnerabilities and Exposures (CVE) These are vulnerabilities in published operating systems and applications software.

The *Common Vulnerability Scoring System (CVSS)* is a system of ranking vulnerabilities that are discovered based on predefined metrics. This system ensures that the most critical vulnerabilities can be easily identified and addressed after a vulnerability test is met. Scores are awarded on a scale of 0 to 10, with the values having the following ranks:

- 0: No issues
- 0.1 to 3.9: Low
- 4.0 to 6.9: Medium
- 7.0 to 8.9: High
- 9.0 to 10.0: Critical

CVSS is composed of three metric groups. These metric groups are described as follows:

- Base includes characteristics of a vulnerability that are constant over time and user environments.
- Temporal includes characteristics of a vulnerability that change over time but not among user environments.
- Environmental includes characteristics of a vulnerability that are relevant and unique to a particular user's environment.

The base metric group includes the following metrics:

- *Access vector (AV)* describes how the attacker would exploit the vulnerability and has three possible values.
  - L stands for Local and means that the attacker must have physical or logical access to the affected system.
  - A stands for Adjacent network and means that the attacker must be on the local network.
  - N stands for Network and means that the attacker can cause the vulnerability from any network.
- *Access complexity (AC)* describes the difficulty of exploiting the vulnerability and has three possible values.
  - H stands for High and means that the vulnerability requires special conditions that are hard to find.
  - M stands for Medium and means that the vulnerability requires somewhat special conditions.
  - L stands for Low and means that the vulnerability does not require special conditions.
- *Authentication (Au)* describes the authentication an attacker would need to get through to exploit the vulnerability and has three possible values.
  - M stands for Multiple and means that the attacker would need to get through two or more authentication mechanisms.
  - S stands for Single and means that the attacker would need to get through one authentication mechanism.
  - N stands for None and means that no authentication mechanisms are in place to stop the exploit of the vulnerability.
- *Availability (A)* describes the disruption that might occur if the vulnerability is exploited and has three possible values.
  - N stands for None and means that there is no availability impact.
  - P stands for Partial and means that system performance is degraded.
  - C stands for Complete and means that the system is completely shut down.
- *Confidentiality (C)* describes the information disclosure that may occur if the vulnerability is exploited and has three possible values.
  - N stands for None and means that there is no confidentiality impact.
  - P stands for Partial and means some access to information would occur.
  - C stands for Complete and means all information on the system could be compromised.

- *Integrity (I)* describes the type of data alteration that might occur and has three possible values.
  - N stands for None and means that there is no integrity impact.
  - P stands for Partial and means some information modification would occur.
  - C stands for Complete and means all information on the system could be compromised.

The CVSS vector will look something like this:

```
CVSS2#AV:L/AC:H/Au:M/C:P/I:N/A:N
```

This vector is read as follows:

`AV:L`

Access Vector: L stands for Local and means that the attacker must have physical or logical access to the affected system.

`AC:H`

Access Complexity: H stands for stands for High and means that the vulnerability requires special conditions that are hard to find.

`Au:M`

Authentication: M stands for Multiple and means that the attacker would need to get through two or more authentication mechanisms.

`C:P`

Confidentiality: P stands for Partial and means some access to information would occur.

`I:N`

Integrity: N stands for None and means that there is no integrity impact.

`A:N`

Availability: N stands for None and means that there is no availability impact.

Control Selection

Once the assets have been classified and their value determined and all vulnerabilities have been identified, controls or mitigations must be selected to address the vulnerabilities. This cannot be done until the level of risk associated with each vulnerability has been determined through one of two methods, qualitative and quantitative risk assessment.

Qualitative Risk Analysis

Qualitative risk analysis does not assign monetary and numeric values to all facets of the risk analysis process. Qualitative risk analysis techniques include intuition, experience, and best-practice techniques, such as brainstorming, focus groups, surveys, questionnaires, meetings, interviews, and Delphi. Although all of these techniques can be used, most organizations will determine the best technique (or techniques) based on the threats to be assessed. Experience and education on the threats are needed.

Each member of the group who has been chosen to participate in the qualitative risk analysis uses their experience to rank the likelihood of each threat and the damage that might result. After each group member ranks the threat possibility, loss potential, and safeguard advantage, data is combined in a report to present to management. All levels of staff should be represented as part of the qualitative risk analysis, but it is vital that some participants in this process should have some expertise in risk analysis.

Quantitative Risk Analysis

A *quantitative risk analysis* assigns monetary and numeric values to all facets of the risk analysis process, including asset value, threat frequency, vulnerability severity, impact, safeguard costs, and so on. Equations are used to determine total and residual risks. The most common equations are for *single loss expectancy (SLE)* and *annual loss expectancy (ALE)*.

The SLE is the monetary impact of each threat occurrence. To determine the SLE, you must know the *asset value (AV)* and the *exposure factor (EF)*. The EF is the percent value or functionality of an asset that will be lost when a threat event occurs. The calculation for obtaining the SLE is as follows:

$$SLE = AV \times EF$$

For example, an organization has a web server farm with an AV of $10,000. If the risk assessment has determined that a power failure is a threat agent for the web server farm and the exposure factor for a power failure is 25 percent, the SLE for this event equals $2,500.

The *annual loss expectancy (ALE)* is the expected risk factor of an annual threat event. To determine the ALE, you must know the SLE and the annualized rate of occurrence (ARO). The ARO is the estimate of how often a given threat might occur annually. The calculation for obtaining the ALE is as follows:

$$ALE = SLE \times ARO$$

Using the previously mentioned example, if the risk assessment has determined that the ARO for the power failure of the web server farm is 50 percent, the ALE for this event equals $1,250.

Cost-Benefit Analysis

Using the ALE, the organization can decide whether to implement controls. If the annual cost of the control to protect the web server farm is more than the ALE, the organization could easily choose to accept the risk by not implementing the control. If the annual cost of the control to protect the web server farm is less than the ALE, the organization should consider implementing the control.

Handling Risk

Risk reduction is the process of altering elements of the organization in response to risk analysis. After an organization understands its total and residual risk, it must

determine how to handle the risk. The following four basic methods are used to handle risk:

Avoidance Terminating the activity that causes a risk or choosing an alternative that is not as risky

Transfer Passing the risk on to a third party, including insurance companies

Mitigation Defining the acceptable risk level the organization can tolerate and reducing the risk to that level

Acceptance Understanding and accepting the level of risk as well as the cost of damages that can occur

Network Topologies

Understanding the types of network topologies that you may see will help you appreciate some of the security measures called for in various scenarios. In this section, you'll learn about some topologies that may exist in your organization.

CAN

The *campus area network* (CAN) comprises the part of the network where data, services, and connectivity to the outside world are provided to those who work in the corporate office or headquarters. It can be further subdivided into the following:

- *Enterprise core* connects the enterprise campus and the intranet data center.
- *Enterprise campus* includes the end devices and provides them access to the outside world and to the intranet data center through the enterprise core.
- *Intranet data center* includes the data center where resources are made available to the enterprise campus and to branch offices though the enterprise core.

Figure 1.3 shows the components of the CAN. It includes two parts that are not part of the enterprise campus (WAN edge and Internet edge) that comprise the networks that are used to connect to the outside world.

FIGURE 1.3 Campus area network

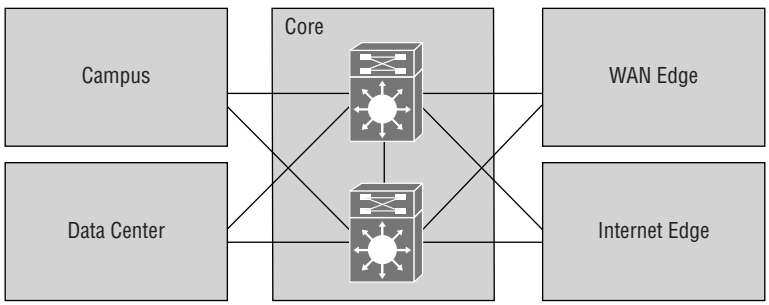

Security issues in the enterprise core include the following:

- Service disruptions (denial of service [DoS], distributed denial of service [DDoS])
- Unauthorized access (intrusions, routing protocol attacks)
- Data leaks and data modifications (packet sniffing, man in the middle [MITM] attacks)

Security issues in the enterprise campus include the following:

- Service disruptions (botnets, malware, DoS)
- Unauthorized access (intrusions, IP spoofing)
- Data leaks and data modifications (packet sniffing, MiTM attacks)
- Identify theft and fraud (phishing, email spam)

Security issues in the intranet data center include the following:

- Unauthorized access (device access, data access, privilege escalation)
- Service disruptions (botnets, DoS)
- Data leaks and data modifications (MITM, malware, scripting, SQL attacks)

WAN

The WAN connection of the organization is called the *enterprise WAN edge* in the Cisco network model. It is one of two modules that are used to connect the CAN to the outside world, the other being the enterprise Internet edge (shown in Figure 1.3). This comprises the provisioned WAN connections to other offices.

Security issues in the enterprise WAN edge include the following:

- Malicious branch client activity (malware, Trojans, botnets)
- Transmission threats (MITM, sniffing)
- Infrastructure attacks (reconnaissance, DoS, service attacks)

Data Center

While the data center may be located in the campus area network, it may also be located in the cloud. The introductions of cloud environments bring many benefits, but they also bring security threats. These threats include the following:

- Account or service hijacking
- Data loss
- Improper device hardening and patching
- DoS attacks
- Insecure APIs and user interfaces
- Malicious provider insiders
- Improper access from other tenants

SOHO

Many of today's workers operate from home rather than in the main office or headquarters. Other users will be operating from smaller branch offices. When this is the case, the *small office/home office (SOHO)* network will connect to the main office via the WAN edge module in cases where the connection is provisioned and via the Internet edge module when the connection leverages the Internet (such as a VPN connection). These two edge modules were shown in Figure 1.3. Since this module interfaces with those two modules, the security issues in the SOHO network will be the same as those present in the Internet edge and WAN edge modules.

Virtual

Today's data centers are increasingly moving to a virtual environment. When a virtual environment is present, it may reside in the campus data center, or it may reside in a cloud data center. Also, it is not unusual to find that the organization has both a physical data center and a virtual data center. Regardless of the exact configuration, there are challenges to securing a virtual environment.

In a virtual environment there are two traffic pathways, one that is used within the virtual environment and one used between the virtual environment and the physical environment. Physical security devices cannot be used to enforce security on the traffic that never leaves a physical host (traffic between VMs located on the same host) or on traffic that never leaves the virtual environment (traffic between VMs on different hosts). The solution is the deployment of virtual security devices such as the Cisco ASAv firewall, the Cisco CSR1000v router, and the Cisco Nexus 1000v switch.

Common Network Security Zones

One of the most basic design principles for a secure network calls for creating *security zones*. These are logical divisions of the network with access controls applied to control traffic between the zones. By organizing resources in these zones and applying the proper access controls, you can reduce the possibility that unauthorized access to data is allowed. In this section, you'll explore four common security zones.

DMZ

A *demilitarized zone (DMZ)* is an area where you can place a public server for access by people you might not trust otherwise. By isolating a server in a DMZ, you can hide or remove access to other areas of your network. You can still access the server using your network, but others aren't able to access further network resources. This can be accomplished using firewalls to isolate your network.

When establishing a DMZ, you assume that the person accessing the resource isn't necessarily someone you would trust with other information. By keeping the rest of the network from being visible to external users, this lowers the threat of intrusion in the internal network.

Any time you want to separate public information from private information, a DMZ is an acceptable option.

The easiest way to create a DMZ is to use a firewall that can transmit in these three directions:

- To the internal network
- To the external world (Internet)
- To the public information you're sharing (the DMZ)

From there, you can decide what traffic goes where; for example, HTTP traffic would be sent to the DMZ, and email would go to the internal network.

Intranet and Extranet

While DMZs are often used to make assets publicly available, extranets are used to make data available to a smaller set of the public—for example, a partner organization. *Intranet* is a term to describe the interior LAN; an *extranet* is a network logically separate from the intranet, the Internet, and the DMZ (if both exist in the design), where resources that will be accessed from the outside world are made available. Access may be granted to customers, business partners, and the public in general. All traffic between this network and the intranet should be closely monitored and securely controlled. Nothing of a sensitive nature should be placed in the extranet.

Public and Private

The purpose of creating security zones such as DMZs is to separate sensitive assets from those that require less protection. Because the goals of security and of performance and ease of use are typically mutually exclusive, not all networks should have the same levels of security.

Information that is of a public nature, or that you otherwise deem not to be of a sensitive nature, can be located in any of the zones you create. However, you should ensure that private corporate data and especially *personally identifiable information (PII)*—information that can be used to identify an employee or customer and perhaps steal their identity—is located only in secure zones and never in the DMZ or the extranet.

VLAN

Network security zones can also be created at layer 2. *Virtual local area networks (VLANs)* are logical subdivisions of a switch that segregate ports from one another as if they were in different LANs. VLANs offer another way to add a layer of separation

between sensitive devices and the rest of the network. For example, if only one device should be able to connect to the finance server, the device and the finance server could be placed in a VLAN separate from the other VLANs. As traffic between VLANs can occur only through a router, ACLs can be used to control the traffic allowed between VLANs.

These VLANs can also span multiple switches, meaning that devices connected to switches in different parts of a network can be placed in the same VLAN regardless of physical location.

Summary

This chapter covered common security principles such as the CIA triad, the goals of which should guide all security initiatives. It also discussed common security terms such as risk, vulnerability, and threat, as well as the proper application of common security zones, such as Intranet, DMZ, and extranets. This chapter also discussed network topologies as seen from the perspective of the Cisco campus area network. Finally, the chapter discussed other methods of network segmentation such as VLANs.

Exam Essentials

Describe the CIA triad. Every security measure you implement should contribute to the achievement of one of three goals. The three fundamentals of security are confidentiality, integrity, and availability (CIA), often referred to as the CIA triad.

Define important security terms. Security professionals should become familiar with terms such as *assets*, *vulnerabilities*, *threats*, *threat agent*, *risk*, *exposure*, and *countermeasures*.

Identify common security zones. Describe *intranet*, *extranet*, *DMZ*, and the *Internet*. Explain their proper use.

Describe common network topologies. Explain various topologies as seen from the perspective of the Cisco campus area network such as the enterprise core, enterprise campus, intranet data center, WAN edge, and intranet edge. Describe the common security issues found in each.

Review Questions

1. Which of the following is *not* one of the CIA triad?
   - **A.** Confidentiality
   - **B.** Integrity
   - **C.** Availability
   - **D.** Accountability
2. Which of the following requires that a user or process is given only the minimum access privilege needed to perform a particular task?
   - **A.** Least privilege
   - **B.** Separation of duties
   - **C.** Job rotation
   - **D.** Mandatory vacation
3. Which of the following occurs when a vulnerability is identified or exploited?
   - **A.** Risk
   - **B.** Threat
   - **C.** Exposure
   - **D.** Countermeasure
4. According to NIST SP 800-30, what is the first step in the risk management process?
   - **A.** Identify threats
   - **B.** Identify impact
   - **C.** Identify vulnerabilities
   - **D.** Identify the assets and their value
5. Which of the following is a measure of how freely data can be handled?
   - **A.** Criticality
   - **B.** Sensitivity
   - **C.** Integrity
   - **D.** Value
6. Which of the following is not a typical commercial data classification level?
   - **A.** Sensitive
   - **B.** Confidential
   - **C.** Secret
   - **D.** Public

7. Which of the following represents data shared only within a meeting in the TLP system?
   - **A.** Amber
   - **B.** White
   - **C.** Red
   - **D.** Green

8. Which of the following is a standard used by the security automation community used to enumerate software flaws and configuration issues?
   - **A.** TLP
   - **B.** CIA
   - **C.** SCAP
   - **D.** CAN

9. Which of the following is *not* a metric group in the Common Vulnerability Scoring System?
   - **A.** Base
   - **B.** Access vector
   - **C.** Temporal
   - **D.** Environmental

10. Which of the following is the monetary impact of each threat occurrence?
    - **A.** ALE
    - **B.** AV
    - **C.** ARO
    - **D.** SLE

11. Which method of handling risk involves defining the acceptable risk level the organization can tolerate and reducing the risk to that level?
    - **A.** Avoidance
    - **B.** Mitigation
    - **C.** Acceptance
    - **D.** Transfer

12. What part of the campus area network includes the end devices and provides them with access to the outside world and to the Intranet data center through the enterprise core?
    - **A.** Intranet data center
    - **B.** Enterprise campus
    - **C.** Enterprise core
    - **D.** Enterprise WAN edge

13. Which of the following is an area where you can place a public server for access by anyone?

A. Intranet

B. DMZ

C. Internet

D. Extranet

14. Which of the following is a logical subdivision of a switch that segregates ports from one another?

A. VLAN

B. VPN

C. DMZ

D. STP

15. Which of the following refers to the data being unaltered by unauthorized individuals?

A. Confidentiality

B. Integrity

C. Availability

D. Accountability

16. Which of the following refers to the practice of using multiple layers of security between data and the resources on which it resides and possible attackers?

A. Default to no access

B. Defense in depth

C. Separation of duties

D. Job rotation

17. Which of the following is the probability that a threat agent will exploit a vulnerability and the impact if the threat is carried out?

A. Risk

B. Threat

C. Exposure

D. Countermeasure

18. Which of the following is a system that uses traffic light colors to classify information assets?

A. DLP

B. VLAN

C. TLP

D. VTP

19. Which component of SCAP refers to vulnerabilities in published operating systems and applications software?

A. CWE

B. CVE

C. CCE

D. CPE

20. Which of the following is the percent value or functionality of an asset that will be lost when a threat event occurs?

A. SLE

B. AV

C. EF

D. ALE