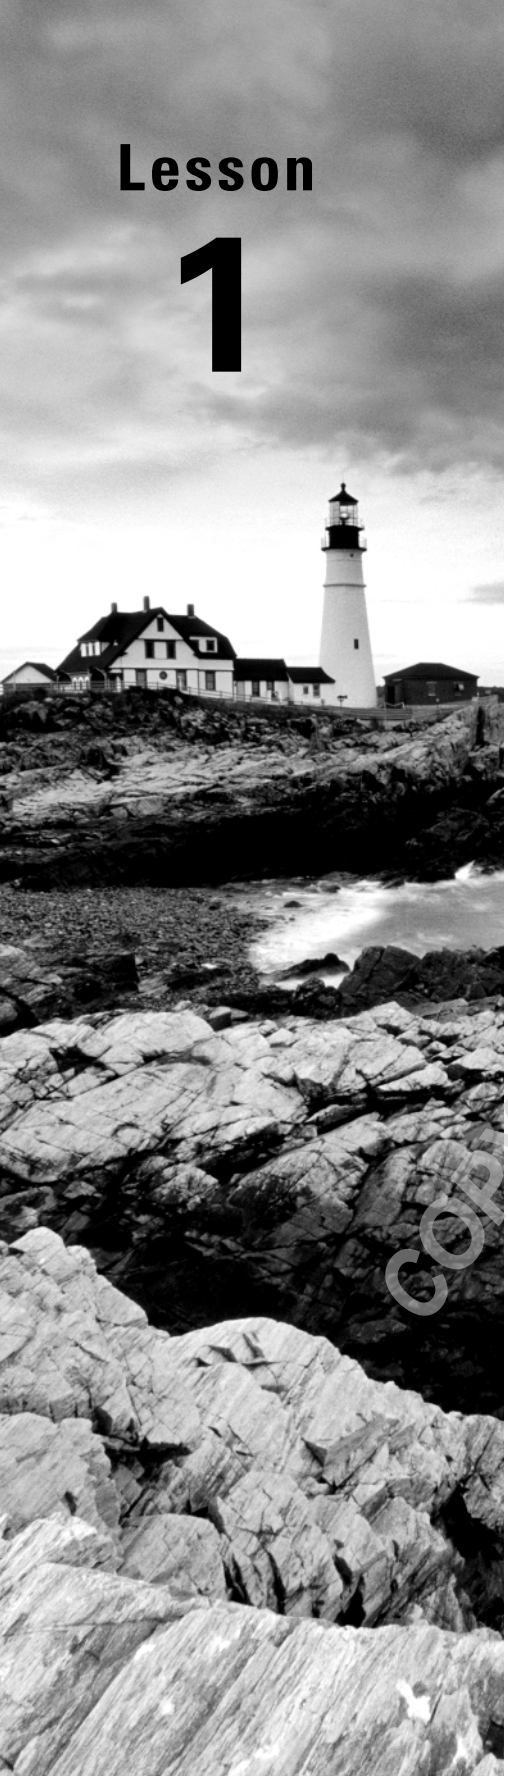

Lesson 1

Understanding Security Layers

Lesson Skill Matrix

Technology Skill	Objective Domain Description	Objective Domain Number
Introducing Core Security Principles	Understand core security principles	1.1
Understanding Physical Security as the First Line of Defense	Understand physical security	1.2
Performing Threat Modeling	Understand core security principles	1.1

Key Terms

access control
attack surface
attack surface analysis
availability
confidentiality
defense in depth
DREAD
egress traffic
flash drive
ingress traffic
integrity
keylogger
mobile devices
Principle of Least Privilege
removable device
residual risk
risk
risk acceptance
risk assessment
risk avoidance
risk mitigation
risk register
risk transfer
separation of duties
social engineering
STRIDE
threat
threat and risk management
threat modeling

Real World Scenario

Lesson 1 Case

When thinking about security, most people start by thinking about their stuff. We all have stuff. We have stuff that we really care about, we have stuff that would be really difficult to replace, and we have stuff that has great sentimental value. We have stuff we really don't want other people to find out about. We even have stuff that we could probably live without. Now think about where you keep your stuff. It could be in your house, your car, your school, your office, in a locker, in a backpack or a suitcase, or a number of other

places. Lastly, think about all of the dangers that could happen to your stuff. People could be robbed or experience a disaster such as a fire, earthquake, or flood. In any case, we all want to protect our possessions no matter where the threat comes from.

At a high level, security is about protecting stuff. In the case of personal stuff, it's about making sure to lock the door when leaving the house, or remembering to take your purse when leaving a restaurant, or even making sure to cover all the presents purchased for Christmas and putting them in the back of the car before heading back into the mall.

Many of the security topics we will discuss in this lesson boil down to the same common sense used every day to protect stuff. In the business environment, the stuff we protect is assets, information, systems, and networks, and we can protect these valuable assets with a variety of tools and techniques that we will discuss at length in this book.

In this lesson, we will start with the basics. We'll look at some of the underlying principles of a security program to set the foundation for understanding the more advanced topics covered later in the book. We'll also discuss the concepts of physical security, which is critical not only for securing physical assets but information assets as well. By the time we're done, you'll have a good idea how to protect stuff for a living.

Introducing Core Security Principles

A fundamental understanding of the standard concepts of security is essential before people can start securing their environment. It's easy to start buying firewalls, but until you understand what needs to be protected, why it needs to be protected, and what it's being protected from, you're just throwing money away.

Certification Ready

List and describe what CIA stands for as it relates to security. Objective 1.1

When working in the security field, one of the first acronyms to be encountered in the information security field is CIA. Not to be confused with the government agency with the same acronym, in information security, this acronym represents the core goals of an information security program. These goals are:

- Confidentiality
- Integrity
- Availability

Understanding Confidentiality

Confidentiality is a concept we deal with frequently in real life. We expect our doctor to keep our medical records confidential. We trust our friends to keep our secrets confidential. In the business world, we define confidentiality as the characteristic of a resource—ensuring access is restricted to only permitted users, applications, or computer systems. What does this mean in reality? Confidentiality deals with keeping information, networks, and systems secure from unauthorized access.

An area where this issue is particularly critical in today's environment is with the high-profile leaking of people's personal information by several large companies. These breaches in confidentiality made the news largely because the information could be used to perpetrate identity theft against the people whose information was breached.

There are several technologies that support confidentiality in an enterprise security implementation. These include the following:

- Strong encryption
- Strong authentication
- Stringent access controls

More Info

Lesson 2 contains more details on these security technologies.

Another key component to consider when discussing confidentiality is how to determine what information is considered confidential. Some common classifications of data are Public, Internal Use Only, Confidential, and Strictly Confidential. The Privileged classification is also used frequently in the legal profession. The military often uses Unclassified, Restricted, Confidential, Secret, and Top Secret. These classifications are then used to determine the appropriate measures needed to protect the information. If information is not classified, there are two options available—protecting all information as if it were confidential (an expensive and daunting task) or treating all information as if it were Public or Internal Use Only and not taking stringent protection measures.

Classify all data and assets—it's the only way to effectively protect them.

Understanding Integrity

We define *integrity* in the information security context as the consistency, accuracy, and validity of data or information. One of the goals of a successful information security program is to ensure that the information is protected against any unauthorized or accidental

changes. The program should include processes and procedures to manage intentional changes, as well as the ability to detect changes.

Some of the processes that can be used to effectively ensure the integrity of information include authentication, authorization, and accounting. For example, rights and permissions could be used to control who can access the information or resource. Also, a hashing function (a mathematical function) can be calculated before and after to show if information has been modified. In addition, an auditing or accounting system can be used that records when changes have been made.

Understanding Availability

Availability is the third core security principle, and it is defined as a characteristic of a resource being accessible to a user, application, or computer system when required. In other words, when a user needs to get to information, it's available to them. Typically, threats to availability come in two types—accidental and deliberate. Accidental threats would include natural disasters like storms, floods, fire, power outages, earthquakes, and so on. This category would also include outages due to equipment failure, software issues, and other unplanned system, network, or user issues. The second category is related to outages that result from the exploitation of a system vulnerability. Some examples of this type of threat would include a denial-of-service attack or a network worm that impacts vulnerable systems and their availability. In some cases, one of the first actions a user needs to take following an outage is to determine into which category an outage fits. Companies handle accidental outages very differently than deliberate ones.

Defining Threat and Risk Management

Threat and risk management is the process of identifying, assessing, and prioritizing threats and risks. A *risk* is generally defined as the probability that an event will occur. In reality, businesses are only concerned about risks that would negatively impact a computing environment. There is a risk that you'll win the lottery on Friday—that's not a risk to actively address, because it would be a positive. A *threat* is a very specific type of risk, and it is defined as an action or occurrence that could result in a breach in the security, outage, or corruption of a system by exploiting known or unknown vulnerabilities. The goal of any risk management plan is to remove risks when possible and to minimize the consequences of risks that cannot be eliminated.

The first step in creating a risk management plan is to conduct a *risk assessment*. Risk assessments are used to identify the risks that might impact an environment.

In a mature risk assessment environment, it is common to record risks in a *risk register*, which provides a formal mechanism for documenting the risks, impacts, controls, and other information required by the risk management program.

After completing an assessment and identifying risks, the next step is to evaluate each risk for two factors. First, determine the likelihood that a risk will occur in the environment. For example, a tornado is much more likely in Oklahoma than in Vermont. A meteor strike is probably not very likely anywhere, although it's the example commonly used to represent the complete loss of a facility when discussing risk. After determining the likelihood of a risk, a user needs to determine the impact of that risk on their environment. A virus on a user's workstation generally has a relatively low impact on the company, although it can have a high impact on the user. A virus on a user's financial system has a much higher impact, although hopefully a lower likelihood.

After evaluating risks, it's time to prioritize them. One of the best mechanisms to assist with the prioritization is to create a risk matrix, which can be used to determine an overall risk ranking. A risk matrix should include the following:

- The risk
- The likelihood that the risk will actually occur
- The impact of the risk
- A total risk score
- The relevant business owner for the risk
- The core security principles that the risk impacts (confidentiality, integrity, and/or availability)
- The appropriate strategy or strategies to deal with the risk

Some additional fields that may prove useful in a risk register include:

- A deliverable date for the risk to be addressed.
- Documentation about the residual risk, which is the risk of an event that remains after measures have been taken to reduce the likelihood or minimize the effect of the event.
- A status on the strategy or strategies to address the risk. These can include status indicators like Planning, Awaiting Approval, Implementation, and Complete.

One easy way to calculate a total risk score is to assign numeric values to the likelihood and impact. For example, rank likelihood and impact on a scale from 1 to 5, where 1 equals low likelihood or low probability, and 5 equals high likelihood or high impact. Then, multiply the likelihood and impact together to generate a total risk score. Sorting from high to low provides an easy method to initially prioritize the risks. Next, review the specific risks to determine the final order in which to address them. At this point, external factors, such as cost or available resources, might affect the priorities.

After prioritizing all risks, there are four generally accepted responses to these risks. These responses include the following:

- Avoid
- Accept
- Mitigate
- Transfer

Risk avoidance is the process of eliminating a risk by choosing to not engage in an action or activity. An example of risk avoidance would be a person who identifies that there is a risk that the value of a stock might drop, so they avoid this risk by not purchasing the stock. A problem with risk avoidance is that there is frequently a reward associated with a risk—avoid the risk and you avoid the reward. If the stock in the example were to triple in price, the risk averse investor would lose out on the reward because he or she wanted to avoid the risk.

Risk acceptance is the act of identifying and then making an informed decision to accept the likelihood and impact of a specific risk. In the stock example, risk acceptance would be the process where a buyer would thoroughly research a company whose stock they are interested in, and after ensuring they are informed, make the decision to accept the risk that the price might drop.

Risk mitigation consists of taking steps to reduce the likelihood or impact of a risk. A common example of risk mitigation is the use of redundant hard drives in a server. There is a risk of hard drive failure in any system. By using redundant drive architecture, users can mitigate the risk of a drive failure by having the redundant drive. The risk still exists, but it has been reduced by a user's actions.

Risk transfer is the act of taking steps to move responsibility for a risk to a third party through insurance or outsourcing. For example, there is a risk that a person may have an accident while driving a car. Purchasing insurance transfers this risk, so that in the event of an accident, the insurance company is responsible to pay the majority of the associated costs.

One other concept in risk management that needs to be covered is *residual risk*. Residual risk is the risk of an event that remains after measures have been taken to reduce the likelihood or minimize the effect of the event. To continue with the car insurance example, the residual risk in the event of an accident would be the deductible a driver has to pay in the event of an accident.

There are many different ways to identify, assess, and prioritize risks. There is no one right way. Use the techniques that best fit the environment and requirements.

While we are discussing risks, we need to look at two final concepts that will help you understand the foundations of security principles and risk management.

Understanding the Principle of Least Privilege

The *Principle of Least Privilege* is a security discipline that requires that a user, system, or application be given no more privilege than necessary to perform its function or job. On its face, this sounds like a very commonsense approach to assigning permissions, and when seen on paper, it is. However, when attempting to try to apply this principle in a complex production environment, it becomes significantly more challenging.

The Principle of Least Privilege has been a staple in the security arena for a number of years, but many organizations struggle to implement it successfully. However, with an increased focus on security from both a business as well as a regulatory perspective, organizations are working harder to build their models around this principle. The regulatory requirements of Sarbanes-Oxley, HIPAA, HITECH, and the large number of state data/privacy breach regulations, coupled with an increased focus by businesses into the security practices of the business partners, vendors, and consultants, are driving organizations to invest in tools, processes, and other resources in order to ensure this principle is followed.

But why is a principle that sounds so simple on paper so difficult to implement in reality? The challenge is largely related to the complexity of a typical environment. It is very easy to visualize how to handle this for a single employee. On a physical basis, they would need access to the building they work in, common areas, and their office.

Logically, the employee needs to be able to log on to their computer, have user access to some centralized applications, access to a file server, a printer, and an internal website. Now, imagine that user multiplied by a thousand. The thousand employees work in six different office locations. Some employees need access to all the locations, while others only need access to their own location. Still others need access to subsets of the six locations; they might need access to the two offices in their region, for example. Some will need access to the data center so they can provide IT support.

Logically, instead of a single set of access requirements, there are multiple departments with varying application requirements. The different user types vary from a user to a power user to an administrator, and you need to determine not only which employee is which type of user, but also manage their access across all the internal applications. Add to this mix new hires, employees being transferred or promoted, and employees who leave the company, and you can start to see how making sure that each employee has the minimum amount of access required to do their job can be a time-intensive activity.

But wait, we're not done. In addition to the physical and user permissions, in many IT environments, applications also have a need to access data and other applications. In order to follow the Principle of Least Privilege, it is important to ensure the applications have the minimum access in order to function properly. This can be extremely difficult when working in a Microsoft Active Directory environment, due to the extremely detailed permissions included in Active Directory. Determining which permissions an application requires to function properly with Active Directory can be challenging in the extreme.

To further complicate matters, in industries where there is heavy regulation, like Finance or Medical, or when regulations like Sarbanes-Oxley are in effect, there are additional requirements that are audited regularly to ensure the successful implementation and validation of privileges across the enterprise.

Getting into a detailed discussion of how to implement and maintain the Principle of Least Privilege is beyond the scope of this book, but there are some high level tools and strategies to be aware of:

Groups Groups can be used to logically group users and applications so that permissions are not applied on a user-by-user basis or application-by-application basis.

Multiple User Accounts for Administrators One of the largest challenges when implementing the Principle of Least Privilege relates to administrators. Administrators are typically also users, and it is seldom a good idea for administrators to perform their daily user tasks as an administrator. To address this issue, many companies will issue their administrators two accounts—one for their role as a user of the company's applications and systems, and the other for their role as an administrator.

Account Standardization The best way to simplify a complex environment is to standardize on a limited number of account types. Each different account type permitted in an environment adds an order of magnitude to the permissions management strategy. Standardizing on a limited set of account types makes managing the environment much easier.

Third-Party Applications There are a variety of third-party tools designed to make managing permissions easier. These can range from account lifecycle management applications to auditing applications and application firewalls.

Processes and Procedures One of the easiest ways to manage permissions in an environment is to have a solid framework of processes and procedures for managing accounts. With this framework to rely on, the support organization doesn't have to address each account as a unique circumstance. They can rely on the defined process to determine how an account is created, classified, permissioned, and maintained.

A perfect implementation of the Principle of Least Privilege is very rare. A best effort is typically what is expected and is achievable.

Understanding Separation of Duties

Separation of duties is a principle that prevents any single person or entity from being able to have full access or complete all the functions of a critical or sensitive process. It is designed to prevent fraud, theft, and errors.

When dealing with orders and payments, it is common to divide those processes into two or more sub-processes. For example, in accounting, the Accounts Receivable employees review and validate bills, and the Accounts Payable employees pay the bills. In any case, those users involved with the critical processes do not have access to the logs. A third set of employees would review and validate what has been occurring and validate that there are no suspicious activities.

When working with IT, while there may be administrators with full access to an application or service, such as a database, the administrators should not be given access to the security logs. Instead, the security administrators regularly review the logs, but these security administrators will not have access to data within the databases. To maintain separation of duties, perform user rights and permissions on a regular basis to ensure that separation of duties is maintained.

Understanding an Attack Surface

One final concept to tackle when discussing core security principles is the idea of an attack surface when evaluating an environment. The concept of an *attack surface* with respect to systems, networks, or applications is another idea that has been around for some time. An attack surface consists of the set of methods and avenues an attacker can use to enter a system and potentially cause damage. The larger the attack surface of an environment, the greater the risk of a successful attack.

In order to determine the attack surface of an environment, it's frequently easiest to divide the evaluation into three components:

- Application
- Network
- Employee

When evaluating the application attack surface, look at things like the following:

- Amount of code in an application
- Number of data inputs to an application
- Number of running services
- Ports on which the application is listening

When evaluating the network attack surface, consider the following:

- Overall network design
- Placement of critical systems
- Placement and rule sets on firewalls
- Other security-related network devices like IDS, VPN, and so on

When evaluating the employee attack surface, consider the following:

- Risk of social engineering
- Potential for human errors
- Risk of malicious behavior

After evaluating these three types of attack surface, you will have a solid understanding of the total attack surface presented by the environment and a good idea of how an attacker might try to compromise the environment.

Performing an Attack Surface Analysis

An attack surface analysis helps to identify the attack surface that an organization may be susceptible to. Because the network infrastructure and necessary services and applications are usually complicated, particularly for medium and large organizations, performing an *attack surface analysis* can also be just as complicated. When completed, the attack surface analysis can be used to determine how to reduce the attack surface.

Certification Ready

What is an attack surface analysis? Objective 1.1

When analyzing a network, the first priority is to determine the security boundaries within an organization. As a minimum, an organization should have an internal network, a DMZ, and the Internet. However, when an organization has multiple sites, or multiple data centers, the organization will also have individual sites, multiple DMZs, and multiple Internet connections. A good place to determine security boundaries is to look at the organization's network documents. Ensure that the organization has proper documentation, which includes network diagrams.

After determining the security boundaries, the next step is to determine everything that connects at those security boundaries. Typically, this includes routers and firewalls, but it might also include some level-3 switches. Next, look at the security mechanisms used for the routers, firewalls, and switches and any security rules associated with those security mechanisms.

With an understanding of the network infrastructure, the next step is to analyze the logs to see which traffic is allowed and which traffic is blocked. *Ingress traffic* is traffic that originates from outside the network's routers and proceeds toward a destination inside the network. *Egress traffic* is network traffic that begins inside a network and proceeds through its routers to its destination somewhere outside of the network.

While network ingress filtering makes Internet traffic traceable to its source, egress filtering helps ensure that unauthorized or malicious traffic never leaves the internal network. Egress traffic might reveal incidents in which an attacker has already gained access to the internal network, or perhaps has gained access to internal users who might be releasing confidential information to the attacker with or without their knowledge. Inter-workload communications should remain internal; they should not transverse the perimeter.

It is important to review egress and ingress traffic on a regular basis. When examining egress and ingress traffic, look at the source and target addresses as well as the ports used. The ports help identify the applications and services to which the traffic packets are related. When creating rules that allow traffic in and out, use descriptive names and consider using templates that can help standardize the setup of multiple firewalls and routers.

In addition to examining egress and ingress traffic, analyze traffic to and from critical systems or systems that contain confidential information. This might help identify problems internally and externally. There might be things that aren't noticed when analyzing egress and ingress traffic.

Testing can also be performed to identify open ports, services, and/or applications that are running on a system and what can be accessed from the outside. There are applications that can test all ports and test for known vulnerabilities. It is also important to configure intrusion detection/prevention systems, including setting alerts that indicate potential threats as they happen.

While analyzing traffic patterns, look at which traffic is encrypted and which traffic is not encrypted. This can help determine which traffic is essential and which traffic could be easily captured. In addition, this helps determine whether encryption should be established for unencrypted data and whether encryption policies need to be established.

To identify application attack services, assess all running network services and applications that communicate with other computers. Then, access best practice guides or hardening guides to learn how to disable any unnecessary program and service so that they cannot be used against you.

When a decision has been made to deploy or adopt a software solution or to build a software solution, it is important to build security into the solution from the beginning. If the organization developed the software solution, make sure the developers and designers are following best practices; their work should be audited from time to time to minimize the risk posed by security vulnerabilities. For third-party applications, choose companies that follow best practices. Ensure they have an update mechanism and process in place for security updates.

Because users often provide the biggest attack surface, remember to review current security policies to make sure that they are being followed. Also, determine whether any policies need to be created or modified. Ensure that all administrators and users are aware of the appropriate policies; if they aren't, ensure that they receive any necessary training.

When evaluating servers, review administrative accounts from time to time to ensure that proper access is provided to the right people. Also, review open sessions. Create and deploy a password policy to make sure that passwords are being changed periodically and that those passwords are strong enough.

Reviewing and reducing attack surfaces should be done periodically to ensure systems are as secure as possible. Also, update the list of attack surfaces as new vulnerabilities are discovered, as new systems are added, and as systems change.

Understanding Social Engineering

One of the key factors to consider when evaluating the employee attack surface is the risk of a social engineering attack. *Social engineering* is a method used to gain access to data, systems, or networks, primarily through misrepresentation. This technique typically relies on the trusting nature of the person being attacked.

In a typical social engineering attack, the attacker will typically try to appear as harmless or respectful as possible. These attacks can be perpetrated in person, through email, or via phone. Attackers will try techniques including pretending to be from a Help Desk or Support Department, claiming to be a new employee, or in some cases even offering credentials that identify them as an employee of the company.

Generally, this attacker will ask a number of questions in an attempt to identify possible avenues to exploit during an attack. If they do not receive sufficient information from one employee, they may reach out to several others until they have sufficient information for the next phase of an attack.

Some techniques for avoiding social engineering attacks include the following:

Be Suspicious Phone calls, emails, or visitors who ask questions about the company, its employees, or other internal information, should be treated with extreme suspicion, and if appropriate, reported to the security organization.

Verify Identity When receiving inquiries that you are unsure of, verify the identity of the requestor. If a caller is asking questions that seem odd, try to get their number so you can call them back. Then, check to ensure that the number is from a legitimate source. If someone approaches with a business card as identification, ask to see a picture ID. Business cards are easy to print, and even easier to take from the "Win a Free Lunch" bowl at a local restaurant.

Be Cautious Do not provide sensitive information unless certain not only of the person's identity but also of the person's right to have the information.

Don't Use Email Email is inherently insecure and prone to a variety of address spoofing techniques. Don't reveal personal or financial information in email. Never respond to email requests for sensitive information and be especially cautious of providing this information after following web links embedded in an email. A common trick is to embed a survey link in an email, possibly offering a prize, or prize drawing, and then asking questions about the computing environment like "How many firewalls do you have deployed?" or "What firewall vendor do you use?" Employees are so accustomed to seeing these types of survey requests in their inbox that they seldom think twice about responding to them.

The key to thwarting a social engineering attack is through employee awareness—if employees know what to look out for, an attacker will find little success.

Linking Cost with Security

When dealing with security, there are some points to keep in mind when developing a security plan. First, security costs money. Typically, the more money is spent, the more secure the information or resources will be (up to a point). So, when examining risk and threats, look at how much the confidential data or resource is worth to the organization if it is compromised or lost and how much money the organization is willing to spend to protect the confidential data or resource.

In addition to cost, strive to make the security seamless to the users who are using or accessing the confidential information or resource. If the security becomes a heavy burden, users will often look for methods to circumvent the security that has been established. Of course, training goes a long way in protecting confidential information and resources, because it will show users what to look for regarding security issues.

Understanding Physical Security as the First Line of Defense

There are a number of factors that need to be considered when designing, implementing, or reviewing physical security measures taken to protect assets, systems, networks, and information. They include understanding site security and computer security, securing removable devices and drives, access control, mobile device security, and identifying and removing keyloggers.

Certification Ready

Why is physical security so important to a server when access to usernames and passwords is needed? Objective 1.2

Most businesses keep some level of control over who accesses their physical environment. There is a tendency when securing computer-related asset and data to only look at the virtual world. Large companies in a location with a data center often use badge readers and/or keypads to provide access to the building and any secure areas. Guards and logbooks are also used to control and track who is in the building. Final layers of security include keys for offices and desk drawers. Similar measures are taken in smaller offices, albeit usually on a smaller scale.

Remember that if someone can get physical access to a server where confidential data is stored, they can, with the right tools and enough time, bypass any security that the server may use to protect the data.

This multi-layered approach to physical security is known as defense-in-depth or a layered security approach. Securing a physical site is more than just putting a lock on the front door and making sure the door is locked. Physical security is a complex challenge for any security professional.

Security does not end with physical security. It is also important to look at protecting confidential information with technology based on authentication, authorization, and accounting including using rights, permissions, and encryption.

Understanding Site Security

Site security is a specialized area of the security discipline. This section introduces some of the more common concepts and technologies that are typically encountered when working in the security field.

Understanding Access Control

Before we jump into site security details, it's important to understand what is meant by access control. *Access control* is a key concept when thinking about physical security. It is also a little confusing, because the phrase is frequently used when discussing information security. In the context of physical security, access control can be defined as the process of restricting access to a resource to only permitted users, applications, or computer systems.

Certification Ready

How does access control relate to site security? Objective 1.1

There are many examples of access control that people encounter every day. These include closing and locking a door, installing a baby gate to keep a toddler from falling down a staircase, and putting a fence around a yard to keep a dog out of the neighbor's flowers.

The difference between the access control practiced in everyday life and the access control encountered in the business world is the nature of what is being protected, and the technologies available to secure them. We will cover these topics in more detail through the rest of this lesson.

FIGURE 1.1 Example of a layered site security model

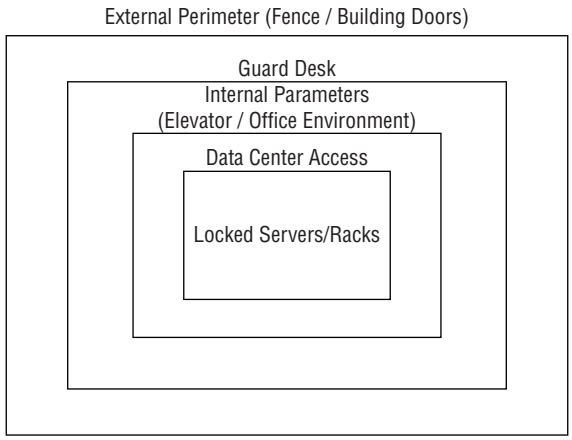

Site security deals with securing the physical premises. One of the fundamental concepts used when designing a security environment is the concept of defense in depth. *Defense in depth* is a concept in which multiple layers of security are used to defend assets. This ensures that if an attacker breaches one layer of defenses, there are additional layers of defense to keep them out of the critical areas of an environment.

A simple example of defense in depth in the "real world" is a hotel room with a locked suitcase. To get into a locked hotel room, a person needs to get the key lock to work. Once they are past the key, there is a deadbolt that must be bypassed. And once they are past the deadbolt, there is still the lock on the suitcase that must be breached.

There are several goals to keep in mind when designing a physical security plan.

Authentication Site security addresses the need to identify and authenticate people permitted access to an area.

Access Control Once a person's identity has been proven and they have been authenticated, site security determines what areas they can access.

Auditing Site security also provides the ability to audit activities within the facility. This can be done through reviewing camera footage, badge reader logs, visitor registration logs, or other mechanisms.

For the purposes of this lesson, we will break the physical premises into three logical areas:

- The external perimeter, which makes up the outermost portion of the location. This typically includes the driveways, parking lots, and any green space the location may support. This does not include things like public roads.
- The internal perimeter, which consists of any buildings on the premises. If the location supports multiple tenants, the internal perimeter is restricted to only the buildings that an employee can occupy.
- Secure areas, which are locations within the building that have additional access restrictions and/or security measures in place. These can include data centers, network rooms, wiring closets, or departments like Research and Development or Human Resources.

Understanding External Perimeter Security

The external security perimeter is the first line of defense surrounding an office. However, security measures in this area probably vary the most of any that we will discuss. When trying to protect a Top Secret government installation, the external perimeter security will consist of multiple fences, roving guard patrols, land mines, and all sorts of other measures that aren't typically used in the corporate world. On the other hand, if an office is in a multi-tenant office park, the external perimeter security may consist of street lights. Most companies fall somewhere in between. Common security measures used for external perimeter security include the following:

- Security cameras
- Parking lot lights
- Perimeter fence
- Gate with guard
- Gate with access badge reader
- Guard patrols

One of the challenges associated with security cameras is that the security camera is only as good as the person monitoring it. Because monitoring cameras is a very expensive, resource-intensive undertaking, in most office environments there will not be anyone actively watching the cameras. Instead, cameras are used after the fact to determine what happened, or who was responsible.

Test an organization's camera playback capabilities regularly. Because cameras are almost always used to review events after the fact, ensure that the system is successfully recording the data.

Understanding the Internal Perimeter

The internal security perimeter starts with the building walls and exterior doors and includes any internal security measures with the exception of any secure areas within the building. Security features that can be used to secure the internal perimeter include the following:

- Locks (exterior doors, internal doors, office doors, desks, filing cabinets, and so on)
- Keypads
- Security cameras
- Badge readers (on doors and elevators)
- Guard desk
- Guard patrols
- Smoke detectors
- Turnstiles
- Mantraps (devices that control access, such as double-doors)

The key security measures implemented in the internal perimeter are utilized to divide the internal space into discrete segments. This is a physical implementation of the Principle of Least Privilege. For example, if the office includes a Finance Department, Human Resources Department, and a Sales Department, it would not be unusual to restrict access to the Finance Department to only people who work in Finance. In general, Human Resources people don't need to be wandering around the Finance area. These segregations may be based on floors, areas, or even a series of offices, depending on the office layout.

Defining Secure Areas

Secure areas would include things like a data center, Research and Development Department, a lab, a telephone closet, a network room, or any other area that requires additional security controls not only from external attackers but also to restrict internal employee access. Secure area security technologies include the following:

- Badge readers
- Keypads

- Biometric technology (fingerprint scanner, retinal scanner, voice recognition, and so on)
- Security doors
- X-ray scanners
- Metal detectors
- Cameras
- Intrusion detection systems (light beam, infrared, microwave, and ultrasonic)

Smaller offices that are not occupied at night may take advantage of remote monitoring, and intrusion detection systems in their internal perimeter. Larger locations typically have some activities going on during nights and weekends, which makes use of these technologies more of a challenge.

Understanding Site Security Processes

While technology forms a significant component when discussing physical security, the processes put in place to support the site security are just as critical. There should be processes at different levels of the site.

In the external perimeter, there might be processes to manage entry to the parking lot through a gate or a process for how often the guards will do a tour of the parking lots. Included in those processes should be how to document findings, how to track entry and exits, and how to respond to incidents. For example, the guard tour process should include instructions on how to handle an unlocked car or a suspicious person or, with the heightened awareness of possible terrorist attacks, how to handle an abandoned package.

In the internal perimeter, processes might include guest sign-in procedures, equipment removal procedures, guard rotation procedures, or details on when the front door is to be left unlocked. In addition, there should probably be processes to handle deliveries, how/when to escort visitors in the facility, and even what types of equipment may be brought into the building. For example, many companies prohibit bringing personal equipment into the office due to the risk that the employee could use their personal laptop to steal valuable company information.

In the secure area layer, there will generally be procedures for controlling who is permitted to enter the data center and how they will access the data center. In addition, you will have multiple mechanisms to ensure that only authorized people are granted access, including locked doors, biometric devices, cameras, and security guards.

Cameras are available on virtually every cell phone on the market today. To ensure that cameras are not used in a facility, plan on taking phones at the door or disabling the camera function.

Understanding Computer Security

Computer security consists of the processes, procedures, policies, and technologies used to protect computer systems. For the purposes of this lesson, computer security will refer to physically securing computers. We will discuss other facets of computer security throughout the rest of the book.

In addition to all the measures we have discussed regarding physical security, there are some additional tools that can be used to secure the actual computers. Before we start discussing the tools, we need to differentiate between the three types of computers we will discuss:

Servers These are computers used to run centralized applications and deliver the applications across a network. This can be an internal network for large businesses or across the Internet for public access. The computer hosting a favorite website is an example of a server. Servers are typically configured with redundant capabilities, ranging from redundant hard drives to fully clustered servers.

Desktop Computers These computers are usually found in office environments, schools, and homes. These computers are meant to be used in a single location and run applications like word processing, spreadsheets, games, and other local applications. They can also be used to interact with centralized applications or browse websites.

Mobile Computers This category includes laptop, notebook, tablet, netbook computers, and smartphones. These are used for the same types of functions as the desktop computer but are meant to be used in multiple locations (for example, home and office). Due to their size, mobile computers are considered to be less powerful than desktop computers, but with the advances in microprocessor technologies and storage technologies, this gap is rapidly narrowing.

When securing a server, the first thing to consider is where the server will be located. Servers are typically significantly more expensive than a desktop or mobile computer and are used to run critical applications, so the types of security typically used with servers are largely location-based. Servers should be secured in data centers or computer rooms, which typically have locked doors, cameras, and other security features we have discussed earlier in the lesson.

If a data center or computer room is not available, other options for securing server computers include the following technologies:

Computer Security Cable A cable that is attached to the computer and to a piece of furniture or wall.

Computer Security Cabinet/Rack A storage container that is secured with a locking door.

Desktop computers are typically secured by the same types of computer security cables that can be used with server computers. Desktop computers are frequently used in secure office environments, or in people's homes, and are not particularly expensive relative to other technologies. Most companies do not take extraordinary measures to protect desktop computers in their offices.

Mobile computers, due to their highly portable nature, have a number of technologies and best practices that can be leveraged to ensure they are not damaged or stolen.

Understanding Mobile Device Security

Mobile devices are one of the largest challenges facing many security professionals today. Mobile devices like laptops, PDAs (Personal Digital Assistants), and smartphones are used to process information, send and receive mail, store enormous amounts of data, surf the Internet, and interact remotely with internal networks and systems. When placing a 32 GB MicroSD memory card in a smartphone that a Senior Vice President can then use to store all the company's Research and Development information, the impact to the company when someone grabs his phone can be staggering. As a result, the security industry makes available a number of technologies for physically securing mobile devices, including the following:

Docking Station Virtually all laptop docking stations are equipped with security features to secure a laptop. This can be with a key, a padlock, or both depending on the vendor and model.

Docking station security only works when the docking station is enabled and secured to an immovable object. It's frequently just as easy to steal a laptop and docking station as it is to just take the laptop.

Laptop Security Cables Used in conjunction with the USS (Universal Security Slot), these cables attach to the laptop and can be wrapped around a secure object like a piece of furniture.

Laptop Safe A steel safe specifically designed to hold a laptop and be secured to a wall or piece of furniture.

Theft Recovery Software An application run on the computer that enables the tracking of a stolen computer so it can be recovered.

Laptop Alarm A motion-sensitive alarm that sounds in the event a laptop is moved. Some are also designed in conjunction with a security cable system so the alarm sounds when the cable is cut. PDAs and smartphones are typically more difficult to secure because they are a newer technology that has exploded in popularity. There are somewhat limited tools available for securing them. For now, configure a password to protect a PDA and phone, enable encryption, and remotely wipe a phone that is managed by an organization. Some of the devices include GPS components that allow users to track a phone or PDA.

Of course, there are some best practices (and, yes, these are based on common sense) that can be followed when securing laptops as well as PDAs or smartphones, including:

Keep your equipment with you. Mobile devices should be kept with you whenever possible. This means keeping mobile devices on your person or in your hand luggage when traveling. Keep mobile devices in sight when going through airport checkpoints.

Use the trunk. When traveling by car, lock the mobile device in the trunk after parking, if you are unable to take the mobile device with you. Do not leave a mobile device in view in an unattended vehicle, even for a short period of time, or left in a vehicle overnight.

Use the safe. When staying in a hotel, lock the mobile device in a safe, if available.

Using Removable Devices and Drives

In addition to mobile devices, another technology that presents unique challenges to security professionals is removable devices and drives. See Figure 1.2 for some examples of common removable devices.

FIGURE 1.2 Some examples of common removable devices

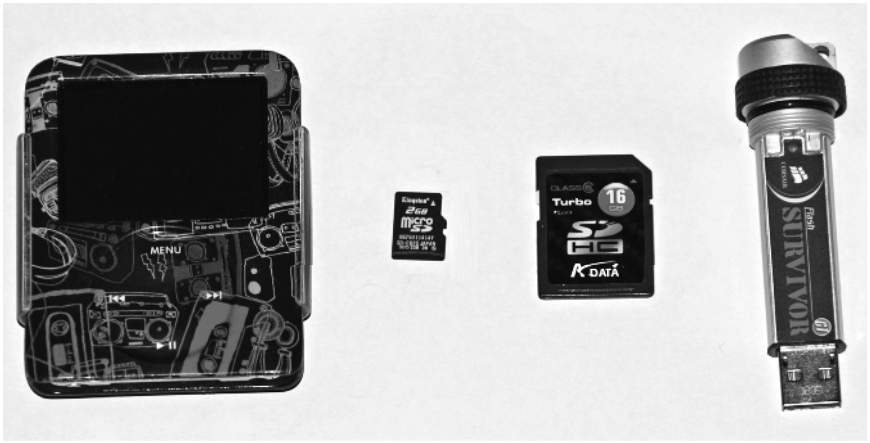

A *removable device* or drive is a storage device that is designed to be removed from the computer without turning the computer off. These devices range from the MicroSD memory card, which is the size of a fingernail and can store 32 GB (or more) of information, to an external hard drive, which can store up to 4 TB of data. CDs, DVDs, and USB drives are also considered removable drives, because they can be used to store critical data and are easily transportable.

These devices typically connect to a computer via a drive or by external communications ports like USB, Firewire, or, in the case of memory cards, through built-in or USB-based readers. These devices are used for a variety of purposes, including backing up critical data, providing supplemental storage, and transferring data between computers. In addition, applications can be run from USB drives. This storage is also used in music players like iPods and Zunes, as well as personal media players like the Archos and Creative's Zen devices.

There are three basic categories of security issues associated with removable storage:

- Loss
- Theft
- Espionage

The loss of the storage device is one of the most common issues people will encounter. USB drives are especially problematic in this regard. Typically the size of a pack of gum or smaller, these drives are often left in conference rooms, hotel rooms, or seat pockets on airplanes. The challenge is how to secure the gigabytes of data that is lost along with these drives. These devices can be protected with authentication and encryption. With Windows 7 and Windows Server 2008 R2, Microsoft released BitLocker To Go, which is used to protect data on mobile storage devices. Some companies may offer their own protection mechanism, such as IronKey. Of course, it is important to impress on users the value of these types of storage. Many users do not give a second thought to throwing a confidential presentation on a *flash drive* (a small drive based on flash memory) for a meeting. As part of the awareness efforts, educate users about the value of data and how easy it is to misplace these portable storage devices.

Theft is a problem with any portable piece of equipment. Many of the same measures discussed with respect to protecting mobile devices apply to these removable storage devices as well. For example, keep drives with you whenever possible. When this is not possible, secure drives in a hotel safe, locked desk drawer, or other secure location. Do not leave portable storage out where it can be easily removed from an accessible area. While the devices themselves are relatively inexpensive, the data on them can be irreplaceable or, worse, confidential.

The final area where these types of devices present a security issue is in conjunction with espionage. Many of these storage devices come in very small form factors, which make them particularly well suited to espionage. Flash drives can be disguised as pens, watches, or even as part of a pocketknife. Even more challenging, a music player or smartphone can include multiple gigabytes of storage. Even if external drives and music players are banned, removing employee's smartphones is virtually impossible. So how do you protect an environment from this type of security threat?

The key to this threat is not to try to defend the environment from the portable devices but instead to protect the data from any unauthorized access. This is where the Principle of Least Privilege is critical—ensure that employees can only access the data, systems, and networks they need to do their jobs so that keeping critical data off portable drives is much easier.

Some environments address the issues associated with removable storage by using hardware or software configurations to prohibit their use. While this can be an effective strategy, it is also an expensive, resource-intensive activity. There are a limited number of businesses where this can be effectively implemented.

More Info

Encryption is frequently used to secure the data on removable drives. This will be discussed in detail in Lesson 2.

Understanding Keyloggers

A *keylogger* is a physical or logical device used to capture keystrokes. An attacker will either place a device between the keyboard and the computer or install a software program to record each keystroke taken and then use software to replay the data to capture critical information like user IDs and passwords, credit card numbers, Social Security numbers, or even confidential emails or other data. There are also wireless keyboard sniffers that can intercept the broadcast keystrokes sent between a wireless keyboard and the computer.

To protect against a physical keylogger, the best tool is visual inspection. Take a look at the connection between the keyboard and the computer. If there is an extra device in between, someone may be trying to capture keystrokes. This is especially important when working with shared or public computers, where attackers will utilize keyloggers to cast a wide net and grab whatever critical data someone might enter.

The best defense against a software keylogger is the use of up-to-date anti-malware software. Many software keyloggers are identified as malware by these applications. User Access Control and host-based firewalls can also be used to prevent a software keylogger from being installed.

To defend against a wireless keyboard sniffer, the best bet is to ensure that a wireless keyboard supports encrypted connections. Most of the current wireless keyboards will either operate in an encrypted mode by default, or at least permit users to configure encryption during installation.

More Info

Lesson 5 contains a more in-depth discussion of anti-malware and workstation firewall technologies.

Performing Threat Modeling

Threat modeling is a procedure for optimizing network security by identifying vulnerabilities, identifying their risks, and defining countermeasures to prevent or mitigate the effects of the threats to the system. It addresses the top threats that have the greatest potential impact to an organization.

Certification Ready

Explain the process of threat modeling. Objective 1.1

Threat modeling is an iterative process; it should be started when designing a system or solution and should be performed throughout the system or solution lifecycle. The reason for multiple passes is that it is impossible to identify all of the possible threats in a single pass. In addition, the infrastructure, system, or solution is always changing, and new threats are found.

The steps to perform threat modeling are:

Identify assets. Identify the valuable assets that the systems must protect.

Create an architecture overview. Gather simple diagrams and related information that show how the systems are connected, both physically and logically. Documentation should include a system, trust boundaries, and data flow.

Decompose the security components and applications. Break down the architecture of the systems and application, including the underlying network and host infrastructure design, security profiles, implementation, as well as the deployment configuration of the systems and applications.

Identify the threats. By examining the current architecture, system, applications, and potential vulnerabilities, identify the threats that could affect the systems and applications.

Document the threats. Document each threat using a common threat template that shows the attributes of each threat.

Rate the threats. Prioritize and address the most significant threats first. The rating process weighs the probability of the threat against the damage that could result should an attack occur. Certain threats might not warrant any action when comparing the risk posed by the threat with the resulting mitigation costs.

One easy way to calculate a total risk score is to assign numeric values to the likelihood and impact. For example, rank likelihood and impact on a scale from 1 to 5, where 1 equals low likelihood or low probability, and 5 equals high likelihood or high impact. Then, multiply the likelihood and impact together to generate a total risk score. Sorting from high to low provides an easy method to initially prioritize the risks. Next, review the specific risks to determine the final order in which to address them. At this point, external factors, such as cost or available resources, might affect the priorities.

STRIDE is an acronym for a threat modeling system that originated at Microsoft. STRIDE is also a mnemonic tool for security threats; it consists of six different categories, as shown in Table 1.1.

TABLE 1.1 STRIDE acronym

Element of STRIDE	Description	Security Properties That Can Reduce STRIDE Risk
Spoofing	Something or someone that pretends to be something that they are not. For example, an attacker could masquerade as a legitimate user or an email can be sent under another domain name or email address.	Authentication

Element of STRIDE	Description	Security Properties That Can Reduce STRIDE Risk
Tampering	Attackers modify or interfere with legitimate data.	Integrity
Repudiation	The user denies performing a certain action, which could be illegal and harmful.	Confirmation
Information Disclosure	A data breach and access to private information occurs, and too much information about a system and its data is accessed by unauthorized individuals.	Confidentiality
Denial-of-service	A service is brought down intentionally or unintentionally resulting in disruptions of applications or services.	Availability
Elevation of Privilege	A user gains privilege access greater than that for which he was approved, potentially accessing restricted data or performing restricted tasks.	Authorization

Use *DREAD* to measure and rank the threats risk level:

Damage Potential How much damage can be inflicted on our system?

Reproducibility Can the attack be reproduced easily?

Exploitability How much effort and experience are necessary?

Affected users If the attack occurs, how many users will be affected?

Discoverability Can the threat be easily discovered?

Rank the threat level on a scale of 0 through 3 or 0 through 10, where the larger the number indicates the greater the threat.

Skill Summary

In this lesson, you learned:

- Before starting to secure an environment, a fundamental understanding of the standard concepts of security is needed.
- CIA (an acronym for Confidentiality, Integrity, and Availability) refers to the core goals of an information security program.
- Confidentiality deals with keeping information, networks, and systems secure from unauthorized access.
- One of the goals of a successful information security program is to ensure integrity or that the information is protected against any unauthorized or accidental changes.

- Availability is defined as a characteristic of a resource being accessible to a user, application, or computer system when required.
- Threat and risk management is the process of identifying, assessing, and prioritizing threats and risks.
- A risk is generally defined as the probability that an event will occur.
- After prioritizing risks, there are four generally accepted responses to these risks: Avoidance, Acceptance, Mitigation, and Transfer.
- The Principle of Least Privilege is a security discipline that requires that a user, system, or application be given no more privilege than necessary to perform its function or job.
- An attack surface consists of the set of methods and avenues an attacker can use to enter a system and potentially cause damage. The larger the attack surface of an environment, the greater the risk of a successful attack.
- The key to thwarting a social engineering attack is through employee awareness. If employees know what to look out for, an attacker will find little success.
- Physical security uses a defense-in-depth or a layered security approach that controls who can physically access resources of an organization.
- Physical premises can be divided into three logical areas: the external perimeter, the internal perimeter, and secure areas.
- Computer security consists of the processes, procedures, policies, and technologies used to protect computer systems.
- Mobile devices and mobile storage devices are one of the largest challenges facing many security professionals today, because of their size and portability.
- A keylogger is a physical or logical device used to capture keystrokes.
- Threat modeling is a procedure for optimizing network security by identifying vulnerabilities, identifying their risks, and defining countermeasures to prevent or mitigate the effects of the threats to the system.

Knowledge Assessment

You can find the answers in the Appendix.

Multiple Choice

1. Which of the following are valid risk responses? (Choose all that apply.)
   - **A.** Mitigation
   - **B.** Transfer
   - **C.** Investment
   - **D.** Avoidance

2. Which of the following are considered removable devices or drives? (Choose all that apply.)
   - **A.** iPod
   - **B.** Netbook
   - **C.** USB flash drive
   - **D.** Burnable DVD drive

3. Which of the following would be considered appropriate security measures for a building's external security perimeter? (Choose all that apply.)
   - **A.** Motion detector
   - **B.** Parking lot lights
   - **C.** Turnstile
   - **D.** Guard patrols

4. When traveling on business and headed out to dinner with a client, which of the following should be done to secure a laptop? (Choose the best answer.)
   - **A.** Lock it in the car trunk.
   - **B.** Store it out of sight in a dresser drawer.
   - **C.** Secure it to a piece of furniture with a laptop security cable.
   - **D.** Check it at the Front Desk.

5. Which of the following refers to the process of eliminating a risk by choosing to not engage in an action or activity?
   - **A.** Mitigation
   - **B.** Residual risk
   - **C.** Avoidance
   - **D.** Acceptance

6. Which of the following technologies could be used to help ensure the confidentiality of proprietary manufacturing techniques for an auto parts manufacturing business? (Choose all that apply.)
   - **A.** Strong encryption
   - **B.** Guard patrols
   - **C.** A laptop safe
   - **D.** Strong authentication

7. The information security acronym CIA stands for which of the following?
   - **A.** Confidentiality, Identity, Access Control
   - **B.** Confidentiality, Integrity, Access Control
   - **C.** Confidentiality, Integrity, Availability
   - **D.** Control, Identity, Access Control

8. Which of the following statements best describes the concept of core security principles?
   - **A.** Core security principles refer to the internal security perimeter when setting up a layered physical security environment.
   - **B.** Core security principles refer to the principles of confidentiality, availability, and integrity.
   - **C.** Core security principles refer to leveraging security best practices.
   - **D.** Core security principles refer to the four methods of addressing risk.

9. As the Chief Security Officer for a small medical records processing company, you have just finished setting up the physical security for your new office. You have made sure that the parking lot is illuminated, that you have guards at the door as well as doing periodic patrols, and you have badge readers throughout the building at key locations. You also have put biometric access technology on the data center door. And of course, you have cameras in the parking lot, building entrances, and the data center entrances.

   This type of implementation is known as: (Choose the best answer)
   - **A.** Access Control
   - **B.** Core Security Principles
   - **C.** Security best practices
   - **D.** Defense in depth

10. Which of the following refers to the process of disabling unneeded services and ports to make the system more secure?
    - **A.** Reducing the attack surface area
    - **B.** Mitigating a Trojan horse
    - **C.** Security avoidance
    - **D.** Defense in depth

11. Which type of network traffic originates from outside the network routers and proceeds toward a destination inside the network?

- **A.** Ingress
- **B.** Egress
- **C.** Traverse
- **D.** Encrypted

Fill in the Blank

1. _____ is characteristic of a business resource—ensuring access is restricted to only permitted users, applications, or computer systems.

2. If a user is deploying technologies to restrict access to a resource, they are practicing the _____ security principle.

3. Deploying multiple layers of security technology to defend assets is called _____.

4. An action or occurrence that could result in a breach in the security, outage, or corruption of a system by exploiting known or unknown vulnerabilities is a(n) _____.

5. A Risk Manager for a medium-sized pharmaceutical company who is asked to perform a formal risk analysis would most likely record the results of the risk assessment in a(n) risk _____.

6. _____ is a method used to gain access to data, systems, or networks, primarily through misrepresentation.

7. The consistency, accuracy, and validity of data or information is called _____.

8. A business traveler notices that there is an extra connector between the keyboard and the computer in a business center. She has most likely encountered a(n) _____.

9. _____ refers to the risk of an event that remains after measures have been taken to reduce the likelihood or minimize the effect of the event.

10. Implementing security measures must always be balanced with _____.

Matching and Identification

What is STRIDE short for?

S _____

T _____

R _____

I _____

D _____

E _____

Build List

Specify the correct order of steps necessary for performing threat modeling.

_____ Create an architecture overview.

_____ Identify assets.

_____ Rate the threats.

_____ Decompose the security components and applications.

_____ Identify the threats.

_____ Document the threats.

Business Case Scenarios

Scenario 1-1: Designing a Physical Security Solution

As the Security Manager for a medium-sized bank, you have been asked to design a security solution to keep a bank robber out of the bank after hours. The three areas of the bank that need to be secured are the parking lot, the building perimeter, and the vault. List what technologies should be used in each area of the bank.

Scenario 1-2: Securing a Mobile Device

An IT Manager for a Legal Services company with 5,000 employees is in the process of rolling out new mobile devices to the Sales Department. Which technologies and best practices should be used to keep these systems physically secure?

Scenario 1-3: Understanding Confidentiality, Integrity, and Availability

A server called Server1 is running Windows Server 2016. On Server1, a folder called Data is created and shared on the C drive. Within the Data folder, subfolders are created with each user's name within the organization. Each person's electronic paycheck is placed in each user's folder. Later, you find out that John was able to go in and change some of the electronic paycheck amounts, while also deleting some of the electronic paychecks. Explain which one (or more) of the CIA components was not followed.

Scenario 1-4: Managing Social Engineering

Your manager at the Contoso Corporation wants to put a training class together for end user security. He wants you to research the Internet for three cases or instances where someone used social engineering to break into a system and describe how they attempted to get access.

Real World Scenario

Workplace Ready: Understanding the Basics

Understanding security concepts is only the first step in learning about security. As a network administrator or security officer, you will be amazed how much going back to the basics will help you plan, implement, and update security procedures.