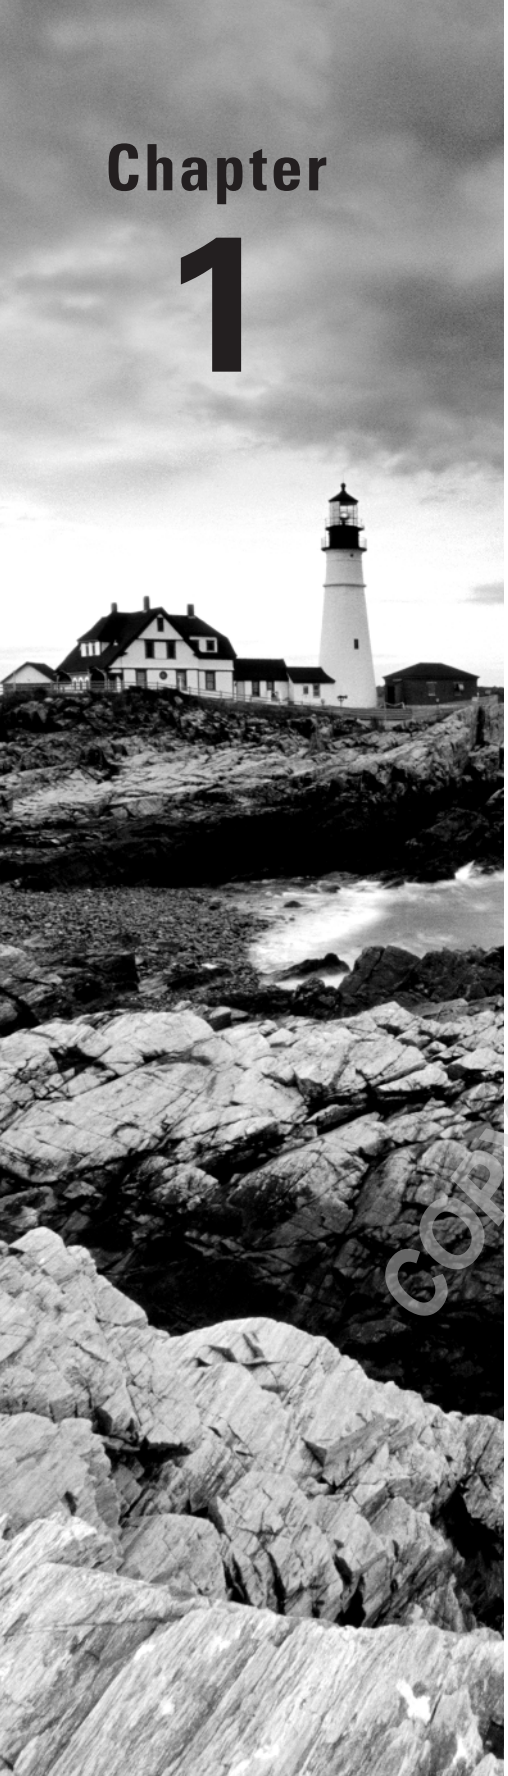

Chapter 1

Managing Risk

THE FOLLOWING COMPTIA SECURITY+ EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **3.8 Explain how resiliency and automation strategies reduce risk.**

- Automation/Scripting: Automated courses of action; Continuous monitoring; Configuration validation
- Templates
- Master image
- Non-persistence: Snapshots; Revert to known state; Rollback to known configuration; Live boot media
- Elasticity
- Scalability
- Distributive allocation
- Redundancy
- Fault tolerance
- High availability
- RAID

✓ **5.1 Explain the importance of policies, plans, and procedures related to organizational security.**

- Standard operating procedure
- Agreement types: BPA; SLA; ISA; MOU/MOA
- Personnel management: Mandatory vacations; Job rotation; Separation of duties; Clean desk; Background checks; Exit interviews; Role-based awareness training (Data owner; System administrator; System owner; User; Privileged user; Executive user); NDA, Onboarding; Continuing education; Acceptable use policy/rules of behavior; Adverse actions
- General security policies: Social media networks/applications; Personal email

✓ **5.2 Summarize business impact analysis concepts.**

- RTO/RPO
- MTBF
- MTTR
- Mission-essential functions
- Identification of critical systems
- Single point of failure
- Impact: Life; Property; Safety; Finance; Reputation
- Privacy impact assessment
- Privacy threshold assessment

✓ **5.3 Explain risk management processes and concepts.**

- Threat assessment: Environmental; Manmade; Internal vs. External
- Risk assessment: SLE; ALE; ARO; Asset value; Risk register; Likelihood of occurrence; Supply chain assessment; Impact; Quantitative; Qualitative; Testing (Penetration testing authorization; Vulnerability testing authorization); Risk response techniques (Accept, Transfer, Avoid, Mitigate)
- Change management

As an administrator, you are responsible. You are responsible for data that gets created, stored, transmitted, viewed, modified, deleted, and just about everything else that can be done with it. Because of this, not only must you enable it to exist, but you must protect it, authenticate it, secure it, and keep it in the form that complies with every applicable law, policy, and regulation. Counter to this are all of the dangers that can befall the data: it can be accidentally deleted, overwritten, stolen, and lost. These potential harms represent *risks*, and you must know the risks involved in working with data. You have to know and accept that data can be corrupted, it can be accessed by those who shouldn't see it, values can be changed, and so on.

If you think that being armed with this knowledge is enough to drive you into taking the steps necessary to keep any harm from happening, however, you are sadly mistaken. One of the actions that administrators can be instructed to take by upper management regarding potential threats is to accept that they exist. If the cost of preventing a particular risk from becoming a reality exceeds the value of the harm that could occur, then a cost-benefit risk calculation dictates that the risk should stand.

Risk calculations weigh a potential *threat* against the *likelihood* or *probability* of it occurring. As frustrating as it may seem, you should accept the fact that some risks, often called *residual risk*, will and must remain. This chapter focuses on risk and the various ways of dealing with it, all of which you will need to understand fully in order to succeed on the Security+ exam.

We will start out by looking at some of the *vernacular*, or terms associated with the field of risk. Then we will move into risk assessment; policies, standards, and guidelines; and change management.

Risk Terminology

Every field of study has a few terms or words that are unique to that particular field in order to help those in the profession to communicate among themselves. The study of risk is no different. A number of terms are associated with risk that will appear at various places in this chapter and throughout the book. The following terms (also found in the online glossary) are those that CompTIA is fond of using and testing. They are provided in order to make it easier for you to know what each is intended to convey.

Security+ Terminology

acceptable use policy/rules of behavior Agreed-upon principles set forth by a company to govern how the employees of that company may use resources such as computers and Internet access.

annual loss expectancy (ALE) A calculation used to identify risks and calculate the expected loss each year.

annualized rate of occurrence (ARO) A calculation of how often a threat will occur. For example, a threat that occurs once every five years has an annualized rate of occurrence of 1/5, or 0.2.

asset value (AV) The assessed value of an item (server, property, and so on) associated with cash flow.

business impact analysis (BIA) A study of the possible impact if a disruption to a business's vital resources were to occur.

business partners agreement (BPA) An agreement between partners in a business that outlines their responsibilities, obligations, and sharing of profits and losses.

exposure factor (EF) The potential percentage of loss to an asset if a threat is realized.

interconnection security agreement (ISA) As defined by NIST (in Publication 800-47), it is "an agreement established between the organizations that own and operate connected IT systems to document the technical requirements of the interconnection. The ISA also supports a Memorandum of Understanding or Agreement (MOU/A) between the organizations."

maximum tolerable downtime (MTD) The maximum period of time that a business process can be down before the survival of the organization is at risk.

mean time between failures (MTBF) The measurement of the anticipated lifetime of a system or component.

mean time to failure (MTTF) The measurement of the average of how long it takes a system or component to fail.

mean time to restore (MTTR) The measurement of how long it takes to repair a system or component once a failure occurs.

memorandum of understanding (MOU)/memorandum of agreement (MOA) Most commonly known as an MOU rather than MOA, this is a document between two or more parties defining their respective responsibilities in accomplishing a particular goal or mission, such as securing a system.

recovery point objective (RPO) The point last known good data prior to an outage that is used to recover systems.

recovery time objective (RTO) The maximum amount of time that a process or service is allowed to be down and the consequences still to be considered acceptable.

Redundant Array of Independent Disks (RAID) A configuration of multiple hard disks used to provide fault tolerance should a disk fail. Different levels of RAID exist.

risk The probability that a particular threat will occur, either accidentally or intentionally, leaving a system vulnerable and the impact of this occurring.

risk acceptance A strategy of dealing with risk in which it is decided the best approach is simply to accept the consequences should the threat happen.

risk analysis An evaluation of each risk that can be identified. Each risk should be outlined, described, and evaluated on the likelihood of it occurring.

risk assessment An evaluation of the possibility of a threat or vulnerability existing. An assessment must be performed before any other actions—such as how much to spend on security in terms of dollars and manpower—can be decided.

risk avoidance A strategy of dealing with risk in which it is decided that the best approach is to avoid the risk.

risk calculation The process of calculating the risks that exist in terms of costs, number, frequency, and so forth.

risk deterrence A strategy of dealing with risk in which it is decided that the best approach is to discourage potential attackers from engaging in the behavior that leads to the risk.

risk mitigation A strategy of dealing with risk in which it is decided that the best approach is to lessen the risk.

risk transference A strategy of dealing with risk in which it is decided that the best approach is to offload some of the risk through insurance, third-party contracts, and/or shared responsibility.

service-level agreement (SLA) An agreement that specifies performance requirements for a vendor. This agreement may use mean time before failure (MTBF) and mean time to repair (MTTR) as performance measures in the SLA.

single loss expectancy (SLE) The cost of a single loss when it occurs. This loss can be a critical failure, or it can be the result of an attack.

single point of failure (SPOF) A single weakness that is capable of bringing an entire system down

vulnerability A flaw or weakness in some part of a system's security procedures, design, implementation, or internal controls that could expose it to danger (accidental or intentional) and result in a violation of the security policy.

Threat Assessment

To protect your resources, you need to be able to identify what threats to them exist—the more specific you can be, the better. It is easy to say that you "might lose data," but that is a danger as opposed to a threat. The *threat* is what would cause you to lose this data. In general terms, a *threat* is anything that can harm your resources, and there are three primary categories of threats that need to be identified and examined:

Environmental Threats from the environment include things such as floods, tornados, hurricanes, and so on. If you share a building with another organization, what would happen if a fire alarm went off in their area? Would sprinklers throughout the entire building be activated and your server room flooded?

Manmade There can be overlap between the categories, and the environmental flooding of a server room could be manmade in nature, caused by an individual holding a match to the bathroom smoke detector.

Internal vs. External If the threat is an individual who is employed by your organization, then it is considered an internal threat. If the individual is not currently employed by your organization, then it is considered an external threat.

One graphical tool that is often used to identify threats is a *risk register*, which is essentially a scatterplot of possible problem areas. With the categories of threats now identified, we will factor them into an assessment of risk in the following sections.

Risk Assessment

Risk assessment is also known as *risk analysis* or *risk calculation*. For purposes of uniformity, we will use *risk assessment* as the term of choice for this discussion. *Risk assessment* deals with the threats, vulnerabilities, and impacts of a loss of information-processing capabilities or a loss of information itself. In simple terms, a *vulnerability* is a weakness that could be exploited by a threat. Each risk that can be identified should be outlined, described, and evaluated for the likelihood of it occurring. The key here is to think outside the box. Conventional threats and risks are often too limited when considering risk assessment.

The chief components of a risk assessment process are outlined here:

Risks to Which the Organization Is Exposed This component allows you to develop scenarios that can help you evaluate how to deal with these types of risks if they occur. An operating system, server, or application may have known risks in certain environments. You should create a plan for how your organization will best deal with these risks and the best way for it to respond to them.

Risks That Need Addressing The risk assessment component also allows an organization to provide a reality check on which risks are real and which are unlikely. This process helps an organization focus on its resources as well as on the risks that are most likely to occur. For example, industrial espionage and theft are likely, but the risk of a hurricane damaging the server room in Indiana is very low. Therefore, more resources should be allocated to prevent espionage or theft as opposed to the latter possibility.

Coordination with BIA The risk assessment component, in conjunction with the *business impact analysis (BIA)*, provides an organization with an accurate picture of the situation facing it. It allows an organization to make intelligent decisions about how to respond to various scenarios.

Real World Scenario

Conducting a Risk Assessment

You've been asked to do a quick assessment of the risks your company faces from a security perspective. What steps might you take to develop an overview of your company's problems?

1. Interview the department heads and the data owners to determine what information they believe requires additional security and to identify the existing vulnerabilities from their perspective.
2. Evaluate the network infrastructure to determine known vulnerabilities and how you might counter them.
3. Perform a physical assessment of the facility to evaluate what physical risks must be countered.

Armed with this information, you have a place to start, and you can determine which countermeasures may be appropriate for the company to mitigate risk.

Computing Risk Assessment

When you're doing a risk assessment, one of the most important things to do is to prioritize. Not everything should be weighed evenly, because some events have a greater likelihood of happening. In addition, a company can accept some risks, whereas others would be catastrophic for the company.

One document you should read is the National Institute of Standards and Technology (NIST) Guide for Conducting Risk Assessments, Publication 800-30. Revision 1 of this document can be found here:

`http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf`

It is worth noting that the revision to the original document refocuses it from being primarily just about risk management to one that strongly emphasizes risk assessment.

Risk Calculations

For purposes of risk assessment, both in the real world and for the exam, you should familiarize yourself with a number of terms to determine the *impact* an event could have:

- *ALE* is the *annual loss expectancy* value. This is a monetary measure of how much loss you could expect in a year.
- *SLE* is another monetary value, and it represents how much you could expect to lose at any one time: the *single loss expectancy*. SLE can be divided into two components:
  - *AV (asset value)*: the value of the item
  - *EF (exposure factor)*: the percentage of it threatened
- *ARO* is the likelihood, often drawn from historical data, of an event occurring within a year: the *annualized rate of occurrence*.

When you compute risk assessment, remember this formula:

$$SLE \times ARO = ALE$$

As an example, if you can reasonably expect that every SLE, which is equal to asset value (AV) times exposure factor (EF), will be the equivalent of $1,000 and that there will be seven such occurrences a year (ARO), then the ALE is $7,000. Conversely, if there is only a 10 percent chance of an event occurring within a year time period (ARO = 0.1), then the ALE drops to $100.

In Exercise 1.1, we'll walk through some risk assessment computations.

EXERCISE 1.1

Risk Assessment Computations

As a security professional, you should know how to compute SLE, ALE, and ARO. Given any two of the numbers, it's possible to calculate the third. Here are three scenarios detailing a hypothetical risk assessment situation followed by the details for figuring out the ALE. They are intended to give you experience working with scenarios similar to those that you may find on the Security+ exam. For this exercise, compute the missing values:

1. You're the administrator of a web server that generates $25,000 per hour in revenue. The probability of the web server failing during the year is estimated to be 25 percent.

A failure would lead to three hours of downtime and cost $5,000 in components to correct. What is the ALE?

The SLE is $80,000 ($25,000 × 3 hours + $5,000), and the ARO is 0.25. Therefore, the ALE is $20,000 ($80,000 × 0.25).

2. You're the administrator for a research firm that works on only one project at a time and collects data through the web to a single server. The value of each research project is approximately $100,000. At any given time, an intruder could commandeer no more than 90 percent of the data. The industry average for ARO is 0.33. What is the ALE?

   The SLE equals $90,000 ($100,000 × 0.9), and the ARO is 0.33. Therefore, the ALE is $29,700 ($90,000 × 0.33).

3. You work at the help desk for a small company. One of the most common requests to which you must respond is to help retrieve a file that has been accidentally deleted by a user. On average, this happens once a week. If the user creates the file and then deletes it on the server (about 60 percent of the incidents), then it can be restored in moments from the shadow copy and there is rarely any data lost. If the user creates the file on their workstation and then deletes it (about 40 percent of the incidents), and if it can't be recovered and it takes the user an average of two hours to re-create it at $12 an hour, what is the ALE?

   The SLE is $24 ($12 × 2), and the ARO is 20.8 (52 weeks × 0.4). Therefore, the ALE equals $499.20 ($24 × 20.8).

Key to any risk assessment is identifying both assets and threats. You first have to identify what it is that you want to protect and then what possible harm could come to those assets. You then analyze the risks in terms of either cost or severity.

Quantitative vs. Qualitative Risk Assessment

Risk assessment can be either *qualitative* (opinion-based and subjective) or *quantitative* (cost-based and objective), depending on whether you are focusing on dollar amounts or simply downtime. The formulas for single loss expectancy (SLE), annual loss expectancy (ALE), and annualized rate of occurrence (ARO) are all based on doing assessments that lead to dollar amounts and are thus quantitative.

To understand the difference between quantitative and qualitative, it helps to use a simple example. Imagine that you get an emergency call to help a small company that you have never heard from before. It turns out that their one and only server has crashed and that their backups are useless. One of the lost files was the only copy of the company's history. This file detailed the company from the day it began to the present day and had the various iterations of the mission statement as it changed over time. As painful a loss as this file represents to the company's culture, it has nothing to do with filling orders and keeping customers happy, and thus its loss is qualitative in nature.

Another loss was the customer database. This held customer contact information as well as the history of all past orders, charge numbers, and so on. The company cannot function

without this file, and it needs to be re-created by pulling all of the hard copy invoices from storage and re-entering them into the system. This loss can be calculated by the amount of business lost and the amount of time it takes to find/re-enter all of the data, and thus it is a quantitative loss.

Whenever you see the word *quantitative*, think of the goal as determining a dollar amount. Whenever you see the word *qualitative*, think of a best guess or opinion of the loss, including reputation, goodwill, and irreplaceable information; pictures; or data that get you to a subjective loss amount.

Risk Measurements

Make sure that you understand the scope and terms of hardware and *service-level agreement (SLA)*–related terms. Doing so can help avoid frustration and prevent unanticipated disruptions from crippling your organization. The following are key measures with which you should be familiar:

Likelihood The meaning of the word *likelihood* is usually self-explanatory; however, actual values can be assigned to likelihood. The National Institute of Standards and Technology recommends viewing likelihood as a score representing the possibility of threat initiation. In this way, it can be expressed either in qualitative or quantitative terms. Table 1.1 shows an assessment scale for the likelihood of threat event initiation adapted from Appendix G of NIST Publication 800-30.

TABLE 1.1 Likelihood assessment scale

Qualitative values	Semi-quantitative values	Description
Very High	10	Adversary is almost certain to initiate threat event.
High	8	Adversary is highly likely to initiate threat event.
Moderate	5	Adversary is somewhat likely to initiate threat event.
Low	2	Adversary is unlikely to initiate threat event.
Very Low	0	Adversary is highly unlikely to initiate threat event.

Guide for Conducting Risk Assessments, National Institute of Standards and Technology, Publication 800-30.

A *supply chain assessment* is similarly used to look at the vendors your organization works with strategically and the potential risks they introduce.

Threat Vectors The term *threat vector* is the way in which an attacker poses a threat. This can be a particular tool that they can use against you (a vulnerability scanner, for example) or the path(s) of attack that they follow. Under that broad definition, a threat vector can be anything from a fake email that lures you into clicking a link (phishing) or an unsecured hotspot (rouge access point) and everything in between.

Mean Time Between Failures The *mean time between failures (MTBF)* is the measure of the anticipated incidence of failure for a system or component. This measurement determines the component's anticipated lifetime. If the MTBF of a cooling system is one year, you can anticipate that the system will last for a one-year period; this means that you should be prepared to replace or rebuild the system once a year. If the system lasts longer than the MTBF, it's a bonus for your organization. MTBF is helpful in evaluating a system's reliability and life expectancy.

Mean Time to Failure Similar to MTBF, the *mean time to failure (MTTF)* is the average time to failure for a nonrepairable system. If the system can be repaired, the MTBF is the measurement to focus on, but if it cannot, then MTTF is the number to examine. Sometimes, MTTF is improperly used in place of MTBF, but as an administrator you should know the difference between them and when to use one measurement or the other.

Mean Time to Restore The *mean time to restore (MTTR)* is the measurement of how long it takes to repair a system or component once a failure occurs. (This is often also referenced as *mean time to repair*.) In the case of a computer system, if the MTTR is 24 hours, this tells you that it will typically take 24 hours to repair it when it breaks.

Although MTTR is considered a common measure of maintainability, be careful when evaluating it because it doesn't typically include the time needed to acquire a component and have it shipped to your location. This author (Emmett Dulaney) once worked with a national vendor who thought MTTR meant mean time to respond—that is, a technician would show up on site within the time called for in the contract, but would only then begin to look at the problem and make a list of any needed supplies. Make sure that the contract or service-level agreement spells out exactly what you want.

Recovery Time Objective The *recovery time objective (RTO)* is the maximum amount of time that a process or service is allowed to be down and the consequences still to be considered acceptable. Beyond this time, the break in business continuity is considered to affect the business negatively. The RTO is agreed on during BIA creation.

Recovery Point Objective The *recovery point objective (RPO)* is similar to RTO, but it defines the point at which the system needs to be restored. This could be where the system was two days before it crashed (whip out the old backup tapes) or five minutes before it crashed (requiring complete redundancy). As a general rule, the closer the RPO matches the item of the crash, the more expensive it is to obtain.

Most SLAs that relate to risk management stipulate the definitions of these terms and how they apply to the agreement. You should understand how these terms are used and what they mean to the vendor and to your organization in order to ensure that there is concurrence.

Assessing Privacy

One area of primary importance for administrators today is privacy. Not only are you charged with keeping data accessible, but that accessibility must be limited to certain parties, and those parties seem to change on a regular basis. In healthcare, for example, records may be limited to only a patient and a doctor, but one patient may have a dozen doctors and all of those doctors need to see that data. Other patients will want their spouse to be able to access their records, while still others won't. And so it goes…

Two privacy-related concepts with which you should be familiar are the *privacy impact assessment (PIA)* and *privacy threshold assessment (PTA)*. A PIA is often associated with a business impact analysis, and it identifies the adverse impacts that can be associated with the destruction, corruption, or loss of accountability of data for the organization. The Department of Homeland Security (DHS), for example, uses it to identify and mitigate privacy risks by telling the public what *personally identifiable information (PII)* it collects, why it is collected, and how it is used, accessed, shared, safeguarded, and stored. According to the DHS, a PIA needs to do three things: ensure conformance with applicable legal, regulatory, and policy requirements for privacy; determine risks and effects; and evaluate protections and alternative processes to mitigate potential privacy risks.

A PTA, on the other hand, is more commonly known as an "analysis" rather than an "assessment." This is the compliance tool used in conjunction with the PIA. An example of the form the DHS uses for this purpose can be found at `https://www.dhs.gov/compliance`.

Two types of testing that can help identify risks are *penetration testing* and *vulnerability testing*. They are particularly useful with identifying threats associated with authorization. Both are discussed in more detail in subsequent chapters.

Acting on Your Risk Assessment

Once you've identified and assessed the risks that exist, for the purpose of the exam, you have four possible responses that you can choose to follow:

Risk Avoidance *Risk avoidance* involves identifying a risk and making the decision not to engage any longer in the actions associated with that risk. For example, a company may decide that many risks are associated with email attachments, and it may choose to forbid any email attachments from entering the network. As part of risk avoidance, the company takes steps to remove the risk, chooses to engage in some other activity, or puts a stop to their exposure to the risk. Avoidance should be based on an informed decision that the best course of action is to deviate from what could lead to exposure to the risk. One of the biggest problems with risk avoidance is that you are actually steering clear of activities from which you may benefit. The most effective risk avoidance strategy to avoid computer crime,

for example, would simply be to avoid using computers at all. Not only is that solution impractical, but it would also prevent companies from adding social value (not to mention monetary value) for their stakeholders.

Risk Transference *Risk transference*, contrary to what the name may imply, does not mean that you shift the risk completely to another entity. What you do instead is share some of the burden of the risk with someone else, such as an insurance company. A typical policy would pay you a cash amount if all the steps were in place to reduce risk and your system was still harmed.

The current push is to move many services to the cloud, hosted by a third-party provider. If you do so, you are engaging in a form of risk transference by relying on that third-party provider for uptime, performance, and security measures. Another risk transference possibility involves employing external consultants for assistance with solutions in areas where internal IT is weak and requiring the external consultants to guarantee their work.

Risk Mitigation *Risk mitigation* is accomplished any time you take steps to reduce risk. This category includes installing antivirus software, educating users about possible threats, monitoring network traffic, adding a firewall, and so on. In Microsoft's Security Intelligence Report, Volume 13, the following suggestions for mitigating risk through user awareness training are listed:

- Keep security messages fresh and in circulation.
- Target new employees and current staff members.
- Set goals to ensure that a high percentage of the staff is trained on security best practices.
- Repeat the information to raise awareness.

CompTIA is fond of risk mitigation and confronting it through the use of routine *audits* that address *user rights* and *permission reviews*; *change management*, the structured approach that is followed to secure a company's assets; and *incident management*, the steps followed when events occur (making sure that controls are in place in order to prevent unauthorized access to, and changes of, all IT assets). Policies addressing data loss or theft need to be in place, and technology controls should be enforced.

Data loss prevention (DLP) systems monitor the contents of systems (workstations, servers, and networks) to make sure that key content is not deleted or removed. They also monitor who is using the data (looking for unauthorized access) and transmitting the data. DLP systems share commonality with network intrusion prevention systems. Other approaches include IPSs, firewalls, and similar devices that can help mitigate risk.

Risk Acceptance *Risk acceptance* is often the choice that you must make when the cost of implementing any of the other responses exceeds the value of the harm that would occur if the risk came to fruition. To truly qualify as acceptance, it cannot be a risk where the administrator or manager is unaware of its existence; it has to be an identified risk for which those involved understand the potential cost or damage and agree to accept it. *Risk acceptance* is nothing more than acknowledging that a risk exists and choosing to do

nothing about it. It does not necessarily mean that you will be affected by the risk, but only that you realize that such a possibility exists. Quite often, this is the choice that you make when the cost of implementing any of the other options exceeds the value of any harm that could occur if the risk is realized. Every firm has a different level of *risk tolerance* (sometimes called a *risk appetite*) that they are willing to accept.

Risk strategies need not be thought of as either/or propositions. It is often possible to combine a bit of mitigation with avoidance. You will often try to combine strategies to reduce your exposure as much as possible. You are then left to accept those issues that cannot be addressed otherwise. In the case of the mailbox analogy, the approach of grouping individual boxes together and placing them all in stone combines elements of both mitigation and transference.

Often you can create interesting or memorable examples to help in understanding or memorizing various lists. This works well for the possible risk responses, too.

Imagine that you are a junior administrator for a large IT department, and you believe that one of the older servers should be replaced with a new one. There are no signs of failure now, but you believe it would be prudent to upgrade before anything disastrous happens. The problem, however, is that all spending requires approval from your superior, who is focused on saving the company as much money as possible and, by doing so, hopes to be considered for a promotion. Thus, she does not want anyone coming up with ways to spend money unnecessarily. You know her well enough to realize that if a problem does occur, she will not hesitate to put all the blame on you in order to save her own career. Table 1.2 shows how you would apply each of the possible risk actions to this scenario.

TABLE 1.2 Risk actions for the scenario

Risk action	Application
Risk avoidance	You begin moving services from the older server to other servers and remove the load to avoid the risk of any services being affected by its demise.
Risk transference	You write up the possibility of the server failing along with details of what you think should be done to prevent it, and you submit your findings to your boss while keeping a copy for yourself. If the server does fail, you have proof that you documented this possibility and made the appropriate parties aware of the situation.
Risk mitigation	You write up the possibility of failure and submit it to your boss while also moving crucial services from that server to others.
Risk acceptance	You know the server could fail but hope that it doesn't. You neither write nor submit reports because you don't want to rock the boat and make your boss unhappy with you. With luck, you'll have transferred to another division before the server ever goes down.

Risk transference, mitigation, and avoidance are all proactive solutions that require planning and implementation ahead of time. Risk acceptance, on the other hand, merely adopts a "do nothing" approach. These constitute the four response strategies that CompTIA expects you to know for the risk management portion of the Security+ exam.

Risks Associated with Cloud Computing

The term *cloud computing* has grown in popularity recently, but few agree on what it truly means. For the purpose of the Security+ exam, *cloud computing* means hosting services and data on the Internet instead of hosting it locally. Some examples of this include running office suite applications such as Office 365 or Google Docs from the web instead of having similar applications installed on each workstation; storing data on server space, such as Google Drive, SkyDrive, or Amazon Web Services; and using cloud-based sites such as `Salesforce.com`.

From an exam standpoint, there are three different ways of implementing cloud computing:

Platform as a Service The *Platform as a Service (PaaS)* model is also known as *cloud platform services*. In this model, vendors allow apps to be created and run on their infrastructure. Two well-known models of this implementation are Amazon Web Services and Google Code.

Software as a Service The *Software as a Service (SaaS)* model is the one often thought of when users generically think of cloud computing. In this model, applications are remotely run over the web. The big advantages are that no local hardware is required (other than that needed to obtain web access) and no software applications need to be installed on the machine accessing the site. The best-known model of this type is `Salesforce.com`. Costs are usually computed on a subscription basis.

Infrastructure as a Service The *Infrastructure as a Service (IaaS)* model utilizes virtualization, and clients pay a cloud service provider for resources used. Because of this, the IaaS model closely resembles the traditional utility model used by electric, gas, and water providers. GoGrid is a well-known example of this implementation.

A number of organizations have examined risk-related issues associated with cloud computing. These issues include the following:

Regulatory Compliance Depending on the type and size of your organization, there are any number of regulatory agencies' rules with which you must comply. If your organization is publicly traded, for example, you must adhere to Sarbanes-Oxley's demanding and exacting rules, which can be difficult to do when the data is not located on your servers. Make sure that whoever hosts your data takes privacy and security as seriously as you do.

User Privileges Enforcing user privileges can be fairly taxing. If the user does not have least privileges (addressed later in this chapter), then their escalated privileges could allow them to access data to which they would not otherwise have access and cause harm to it—intentional or not. Be cognizant of the fact that you won't have the same control over user

accounts in the cloud as you do locally, and when someone locks their account by entering the wrong password too many times in a row, you or they could be at the mercy of the hours that the technical staff is available at the provider.

Data Integration/Segregation Just as web-hosting companies usually put more than one company's website on a server in order to be profitable, data-hosting companies can put more than one company's data on a server. To keep this from being problematic, you should use encryption to protect your data. Be aware of the fact that your data is only as safe as the data with which it is integrated. For example, assume that your client database is hosted on a server that another company is also using to test an application that they are creating. If their application obtains root-level access at some point (such as to change passwords) and crashes at that point, then the user running the application could be left with root permissions and conceivably be able to access data on the server for which they are not authorized, such as your client database. Data segregation is crucial; keep your data on secure servers.

Data integration is equally important. Make sure that your data is not comingled beyond your expectations. It is not uncommon in an extranet to pull information from a number of databases in order to create a report. Those databases can be owned by anyone connected to the extranet, and it is important to make certain that the permissions on your databases are set properly to keep other members from accessing more information than you intended to share.

Among the groups focused on cloud security issues, one worth noting is the Cloud Security Alliance (`https://cloudsecurityalliance.org`). They have published a number of whitepapers on security-related issues that can be found on their site and should be considered highly recommended reading for security administrators.

Risks Associated with Virtualization

If cloud computing has grown in popularity, *virtualization* has become the technology du jour. Virtualization consists of allowing one set of hardware to host multiple virtual machines. It is in use at most large corporations, and it is also becoming more common at smaller businesses.

Some of the possible security risks associated with virtualization include the following:

Breaking Out of the Virtual Machine If a disgruntled employee could break out of the virtualization layer and were able to access the other virtual machines, they could access data that they should never be able to access.

Intermingling Network and Security Controls The tools used to administer the virtual machine may not have the same granularity as those used to manage the network. This could lead to privilege escalation and a compromise of security.

Most virtualization-specific threats focus on the hypervisor. *Hypervisor* is the virtual machine monitor—that is, the software that allows the virtual machines to exist. If the hypervisor can be successfully attacked, the attacker can gain root-level access to all virtual systems. Although this is a legitimate issue, and one that has been demonstrated as possible in most systems (including VMware, Xen, and Microsoft Virtual Machine), it is one that has been patched each time it has arisen. The solution to most virtualization threats is always to apply the most recent patches and keep the system(s) up to date. Be sure to look for and implement suggestions that the vendor of your virtualization system may have published in a hardening guide.

Developing Policies, Standards, and Guidelines

The process of implementing and maintaining a secure network must first be addressed from a policies, standards, and guidelines perspective. This sets the tone, provides authority, and gives your efforts the teeth they need to be effective. Policies and guidelines set a standard of expectation in an organization. The process of developing these policies will help everyone in an organization become involved and invested in making security efforts successful. You can think of policies as providing high-level guidance on large issues. Standards tell people what is expected, and guidelines provide specific advice on how to accomplish a given task or activity.

There is a difference between "top-down policies" (those that use the support of upper management) and "bottom-up policies" (often generated by the IT department with little intradepartmental support).

The following sections discuss the policies, standards, and guidelines that you'll need to establish in order for your security efforts to be successful.

Implementing Policies

Policies provide the people in an organization with guidance about their expected behavior. Well-written policies are clear and concise, and they outline the consequences when they aren't followed. A good policy contains several key areas besides the policy itself.

Scope Statement A good policy has a *scope statement* that outlines what the policy intends to accomplish and which documents, laws, and practices the policy addresses. The scope statement provides background to help readers understand what the policy is about and how it applies to them.

The scope statement is always brief—usually not more than a single sentence in length.

Policy Overview Statement A *policy overview statement* provides the goal of the policy, why it's important, and how to comply with it. Ideally, a single paragraph is all you need to provide readers with a sense of the policy.

Policy Statement Once the policy's readers understand its importance, they should be informed about the substance of the policy. A *policy statement* should be as clear and unambiguous as possible. The policy may be presented in paragraph form, as bulleted lists, or as checklists.

The presentation will depend on the policy's target audience as well as its nature. If the policy is intended to help people determine how to lock up the building at the end of the business day, for example, it might be helpful to provide a specific checklist of the steps that need to be taken to accomplish this task.

Accountability Statement This policy should address who (usually expressed as a position, not the actual name of an individual) is responsible for ensuring that the policy is enforced. The *accountability statement* provides additional information to the reader about who to contact if a problem is discovered. It should also indicate the consequences of not complying with the policy.

The accountability statement should be written in such a way as to leave no room for misinterpretation on the part of users.

Exception Statement Sometimes, even the best policy doesn't foresee every eventuality. The *exception statement* provides specific guidance about the procedure or process that must be followed in order to deviate from the policy. This may include an escalation contact in the event that the person who is dealing with the situation needs to know whom to contact next.

The policy development process is often time-consuming. The advantage of this process, though, is that the decisions can be made in advance and can be sent to all involved parties so that the policy doesn't have to be restated over and over again. In fact, formally developing policies saves time and provides structure: Instead of using valuable time trying to figure out what to do, employees will know exactly what to do.

Incorporating Standards

A *standard* deals with specific issues or aspects of a business. Standards are derived from policies. A standard should provide enough detail that an audit can be performed to determine whether the standard is being met. Standards, like policies, have certain structural aspects in common.

The following five points are the key aspects of standards documents:

Scope and Purpose The *standards document* should explain or describe the intention. If a standard is developed for a technical implementation, the scope might include software, updates, add-ins, and any other relevant information that helps the implementer carry out the task.

Roles and Responsibilities This section of the standards document outlines who is responsible for implementing, monitoring, and maintaining the standard. In a system configuration, this section would outline what the customer is supposed to accomplish and what the installer is supposed to accomplish. This doesn't mean that one or the other can't exceed those roles; it means that, in the event of confusion, it's clear who is responsible for accomplishing which tasks.

Reference Documents This section of the standards document explains how the standard relates to the organization's different policies, thereby connecting the standard to the underlying policies that have been put in place. In the event of confusion or uncertainty, it also allows people to go back to the source and figure out what the standard means. You'll encounter many situations throughout your career where you're given a standard that doesn't make sense. Frequently, by referring to the policies, you can figure out why the standard was written as it was. Doing so may help you carry out the standard or inform the people responsible for the standard of a change or problem.

Performance Criteria This part of the standards document outlines how to accomplish the task. It should include relevant baseline and technology standards. Baselines provide a minimum or starting point for the standard. Technology standards provide information about the platforms and technologies. Baseline standards spell out high-level requirements for the standard or technology.

An important aspect of performance criteria is benchmarking. You need to define what will be measured and the metrics that will be used to do so.

If you're responsible for installing a server in a remote location, for example, the standards spell out what type of computer will be used, what operating system will be installed, and any other relevant specifications.

Maintenance and Administrative Requirements These standards outline what is required to manage and administer the systems or networks. For instance, in the case of a physical security requirement, the frequency with which locks or combinations are changed would be addressed.

As you can see, the standards documents provide a mechanism for both new and existing standards to be evaluated for compliance. The process of evaluation is called an *audit*. Increasingly, organizations are being required to conduct regular audits of their standards and policies.

Following Guidelines

Guidelines are slightly different from either policies or standards. *Guidelines* help an organization implement or maintain standards by providing information on how to accomplish the policies and maintain the standards.

Guidelines can be less formal than policies or standards because their nature is to help users comply with policies and standards. An example might be an explanation of how to install a service pack and what steps should be taken before doing so.

Guidelines aren't hard-and-fast rules. They may, however, provide a step-by-step process to accomplish a task. Guidelines, like standards and policies, should contain background information to help a user perform the task.

The following four items represent the minimum contents of a good guidelines document:

Scope and Purpose The *scope and purpose* section provides an overview and statement of the guideline's intent. It is not uncommon to see the heading "Purpose and Scope" or "Scope and Purpose" at the beginning of a document followed by verbiage to the effect: "This document contains the guidelines and procedures for the assignment and use of *xyz* and establishes the minimum requirements for governing the acceptable use of…"

Where scope and purpose are two separate headings, the information beneath the "Purpose" section states why it exists (for example, "This policy establishes guidelines and minimum requirements governing…"), and the "Scope" section tells to whom it applies (for instance, "This policy applies to any employee who…").

Roles and Responsibilities This section of the guidelines identifies which individuals or departments are responsible for accomplishing specific tasks. This may include implementation, support, and administration of a system or service. In a large organization, it's likely that the individuals involved in the process will have different levels of training and expertise. From a security perspective, it could be disastrous if an unqualified technician installed a system without guidelines.

Guideline Statements The *guideline statements* provide the step-by-step instructions or procedures on how to accomplish a task in a specific manner. Again, these are guidelines or recommendations—they may not be hard-and-fast rules and can even include shortcuts and suggestions.

Operational Considerations A procedure's *operational considerations* specify and identify what duties are required and at what intervals. This list might include daily, weekly, and monthly tasks. Guidelines for systems backup, for example, might provide specific guidance as to which files and directories must be backed up and how frequently.

Guidelines help an organization in three ways:

- If a process or set of steps isn't performed routinely, experienced support and security staff will forget how to do them; guidelines will help refresh their memory.
- When you're trying to train someone to do something new, written guidelines can reduce the new person's learning curve.
- When a crisis or high-stress situation occurs, guidelines can keep you from coming unglued.

Business Policies to Implement

Business policies also affect the security of an organization. They address organizational and departmental business issues as opposed to corporate-wide personnel issues. When developing your business policy, you must consider these primary areas of concern. Policies can be divided into two general categories: those for vendors and those for personnel.

Implementing Policies for Vendors

The overriding policy for operations is the *standard operating procedure (SOP)*. This serves as the baseline for business and, if properly written, covers what is expected on a regular basis. More importantly, from a crisis management standpoint, it outlines what to do when things aren't running as well as they should, such as which vendor to call when the communications server crashes, who to notify when their keypads won't allow access to the server room, and so on.

Service-level agreements (SLAs) were discussed previously, and equally important are *business partner agreements (BPAs)*. These outline responsibilities and obligations (as well as the sharing of profits and losses) between business partners. These documents can be important when you need authorization before or after a disaster to get the help and equipment needed to have everything functioning as it should be.

For the most part synonymous, *memorandum of understanding (MOU)* and *memorandum of agreement (MOA)* define the terms and conditions for securely sharing data and information resources. It is important that it identify the purpose for the interconnection's existence. It should also identify who the relevant authorities are within each organization and define their responsibilities. Since nothing lasts forever, conditions of terminating the agreement should be included in it as well as expected cost apportionment.

An *interconnection security agreement (ISA)* documents the technical and security requirements for establishing, operating, and maintaining the interconnection. It works in conjunction with the MOU/A between organizations by spelling out the requirements for connecting the IT systems and describing the security controls to be used to protect the systems and data, and it includes any necessary topological drawings of the interconnection.

Although the two can overlap, the general guideline is that the ISA specifies the technical and security requirements of the interconnection, whereas the MOU/A defines the responsibilities of each organization. As such, the MOU/A should not have technical details, such as how the interconnection is going to be established or maintained within it. Those details are contained within the ISA.

Given the sensitive nature of the information contained within the ISA and MOU/A, it is imperative that their physical copies be stored in a safe and secure location to keep them from prying eyes. It is also important that electronic copies be protected from those same prying eyes.

Implementing Policies for Personnel

For exam purposes, there are more policies related to personnel—employees and contractors—with which you should be familiar than for vendors. The following sections look at them and guidelines associated with them.

Mandatory Vacations A *mandatory vacation policy* requires employees to take time away from work to refresh, and it is primarily used in jobs related to the financial sector. As contradictory as it may seem, an employee who doesn't take their vacation time can be detrimental to their own health as well as the organization's health. If the company becomes too dependent on one person, they can end up in a real bind if something should happen to that person. Not only does mandatory vacation give the employee a chance to refresh, but it also gives the company a chance to make sure that others can fill in any gaps in skills, and it satisfies the need to have replication or duplication at all levels. Mandatory vacations also provide an opportunity to discover fraud.

Job Rotation A *job rotation policy* defines intervals at which employees must rotate through positions. Similar in purpose to mandatory vacations, it helps to ensure that the company does not become too dependent on one person (who then has the ability to do enormous harm). Rotate jobs on a frequent enough basis so that you are not putting yourself—and your data—at the mercy of any one administrator. Just as you want redundancy in hardware, you want redundancy in abilities.

When one person fills in for another, such as for mandatory vacations, it provides an opportunity to see what the person is doing and potentially uncover any fraud.

Separation of Duties Policies *Separation of duties policies* are designed to reduce the risk of fraud and to prevent other losses in an organization. A good policy will require more than one person to accomplish key processes. This may mean that the person who processes an order from a customer isn't the same person who generates the invoice or deals with the billing.

Separation of duties helps prevent various problems, such as an individual embezzling money from a company. To embezzle funds successfully, an individual would need to recruit others to commit an act of *collusion*—that is, an agreement between two or more parties established for the purpose of committing deception or fraud. Collusion, when part of a crime, is also a criminal act in and of itself.

In addition, separation-of-duties policies can help prevent accidents from occurring in an organization. Let's say that you're managing a software development project. You want someone to perform a quality assurance test on a new piece of code before it's put into production. Establishing a clear separation of duties prevents development code from entering production status until quality testing is accomplished.

Many banks and financial institutions require multiple steps and approvals to transfer money. This helps reduce errors and minimizes the likelihood of fraud.

Very small assaults are often called "salami attacks." In banking, various forms of salami attacks can occur, such as shaving a few cents from many accounts, rounding to whole numbers and compiling the remainder into one account, and so on.

Clean Desk As secure as data within a computer system may be, equally insecure are printed copies of the data resting in a pile on someone's desk. A clean desk policy increases the physical security of data by requiring employees to limit what is on their desk to what they are working on at the present time, safely securing (locking in a drawer, for example) anything not presently needed.

Background Checks Given the need to protect data, one of the best ways to do so is to make sure those who are given access to it can be trusted. All potential employees should be thoroughly screened with an extensive background check before being hired and given access to the computer systems.

Nondisclosure Agreement A nondisclosure agreement (NDA) is a legal contract intended to cover confidentiality. An NDA can be crafted and accepted to cover almost any scale. It can, for example, be limited to only what is discussed in a single meeting with a vendor or cover everything you learned during the course of your employment.

Onboarding The process used to train a new employee and bring them up to speed with the organization, its clients, its products, and so forth is known as onboarding. From the organization's prospective, this process can represent a substantial investment in the new employee and thus the organization typically wants to make sure they made a smart hiring decision and are going to keep the employee for a while. From the new employee's standpoint, they are usually much more valuable after the onboarding has concluded.

Continuing Education Lifelong learning is a necessity in the workplace of today. Not only does additional learning add to the skills an employee has and make them more valuable to the employer, but often it is required to maintain certifications. Most CompTIA certifications today (including Security+), for example, require a minimum number of continuing education units (CEUs) in order for the certified administrator to keep that certification over time.

Exit Interviews One of the best ways to find problems is to listen—not talk—to those with whom you work. The last real opportunity for that communication occurs when an individual leaves the organization and they are given an exit interview. Never bypass this learning opportunity, and be sure to listen carefully to what is being said and ask questions that can help you determine if any changes should be made.

Role-Based Awareness Training Not all employees are equal, and that is especially truly when it comes to data access. Depending on your organization, you may use groups with lots of different names and assign users to those groups. You should, however, at a minimum have the following: Data Owner, System Administrator, System Owner, User, Privileged User, and Executive User.

When deciding what group each user should be a member of, the *least privilege policy* should be used, particularly when assigning permissions. Give users only the permissions that they need to do their work and no more. For example, a temporary employee should never have the right to install software, a receptionist does not need the right to make backups, and so on. Every operating system includes the ability to limit users based on groups and individual permissions, and your company should adhere to the policy of always applying only those permissions users need and blocking those that they do not.

Any time you see the phrase "least privilege," always equate it with giving only the minimum permissions needed to do the work that must be done.

Acceptable Use Policies *Acceptable use policies (AUPs)* describe how the employees in an organization can use company systems and resources, both software and hardware. This policy should also outline the consequences for misuse. In addition, the policy (also known as a *use policy*) should address the installation of personal software on company computers and the use of personal hardware such as USB devices. When portable devices are plugged directly into a machine, they bypass the network security measures (such as firewalls) and allow data to be copied in what is known as *pod slurping*. This can also be done if employees start using free cloud drives instead, and that scenario should be addressed in the AUP.

Even secure workstations that do not contain traditional media devices (CD, DVD, and so forth) usually contain USB ports. Unless those ports are disabled, a user can easily connect a flash drive and copy files to and from it. Not only should you make every effort to limit USB ports, but you should also have the use of such devices spelled out in the acceptable use policy to circumvent the "I didn't know" defense.

Real World Scenario

The Trouble with Not Having a Policy

A few years ago, an employee in a large company was using corporate computer systems to run a small accounting firm that he had started. He was using the computers on his own time. When this situation was discovered, he was immediately fired for the misuse of corporate resources. He sued the company for wrongful discharge and won the case. The company was forced to hire him back and pay his back wages, and he was even awarded damages. The primary reason the company lost the case was that its acceptable use policy didn't state that he couldn't use company computers for personal work—only that he couldn't use them for personal work during work hours. The company was unable to prove that he did the personal work during work hours.

Every acceptable use policy today should include a section on smartphone usage (and even presence) within the workplace. Although a smartphone is a convenience for employees (they can now more easily receive and make personal calls at work), it can be a headache for the security administrator. Most smartphones can store files in the same way as any USB device, and they can be used to copy files to and from a workstation. Additionally, the camera feature on most phones makes it possible for a user to take pictures of things such as company documents, servers, and physical security implementation, among many other things that the company may not want to share. For this reason, most secure facilities have stringent restrictions on the presence of smartphones within the facility.

Make sure your acceptable use policies provide your company with adequate coverage regarding all acceptable uses of corporate resources.

Adverse Actions It is a sad but true condition in the workplace that administrative (usually adverse) actions must be taken against employees. This could be due to an employee being forced to take an administrative leave, being terminated, or any number of other situations. An adverse action policy should be in place detailing exactly what must be done—suspending the user's account, revoking privileges, and so forth.

The more detailed the policy, the less opportunity there is for something important to fall through the cracks and put all of your valuable data at risk. You may truly work with the best co-workers on the planet, but it is surprising how vindictive they can become if they feel they've been wronged. You will be much happier if you can stick to preventive measures and head off problems before they ever have the opportunity to manifest themselves.

General Security Policies *Security policies* define what controls are required to implement and maintain the security of systems, users, and networks. This policy should be used as a guide in system implementations and evaluations, addressing what one might consider common sense.

Issues that should be defined in these policies include the difference between company and personal email and how to deal appropriately with social media—reminding users that they are seen as an extension of the company and that they must always put the interests of the company first. Both social media and personal use of email are items that can cause companies to lose big in the marketplace and the courtroom. Users need to understand that with a company email account and/or use of company resources, they are serving as representatives of the company and any offensive comments they make could reflect poorly on the company.

Network/Application Policies

As you just learned, anything owned by the company is thought to represent the company in both the marketplace and the courtroom. This is true of employees using company email accounts for personal email, the use of company computers for social media, and for all other applications and networks that the employees access while at work or in the line of work.

Understanding Control Types and False Positives/Negatives

Risk assessment and analysis involves calculating potential risks and making decisions based on the variables associated with those risks (likelihood, ALE, impact, and so forth). Once you've identified risks that you want to address with actions other than avoidance, you put controls in place to address those risks.

Control Types

The National Institute of Standards and Technology places controls into various types. The control types fall into three categories: Management, Operational, and Technical, as defined in Special Publication 800-12. Table 1.3 lists the control types and the controls with which they are associated.

TABLE 1.3 Control types and controls

Control type	Controls
Management	Risk Assessment
Management	Planning
Management	System and Services Acquisition
Management	Certification, Accreditation, and Security Assessment
Operational	Personnel Security
Operational	Physical and Environmental Protection
Operational	Contingency Planning
Operational	Configuration Management
Operational	Maintenance
Operational	System and Information Integrity
Operational	Media Protection
Operational	Incident Response
Operational	Awareness and Training
Technical	Identification and Authentication

Control type	Controls
Technical	Access Control
Technical	Audit and Accountability
Technical	System and Communication Protection

Another series of security controls worth examining is NIST 800-53, used by government and industry and viewed as more of a global standard. As of this writing, the current publication is revision 4.

Although we discussed risk assessment in this chapter, we address most of the other controls in subsequent chapters.

After you have implemented security controls based on risk, you must perform routine audits. These audits should include reviews of user rights and permissions as well as specific events. Pay particular attention to false positives and negatives.

False positives are events that aren't really incidents. Event flagging is often based on established rules of acceptance (deviations from which are known as *anomalies*) and things such as attack signatures. If the rules aren't set up properly, normal traffic may set off an analyzer and generate an event. You don't want to declare an emergency unless you're sure that you have one. The opposite of a false positive is a *false negative*. With a false negative, you are not alerted to a situation when you should be alerted. In this case, you miss something crucial and it slips right by.

Error Types

A number of error types exist beyond what you need to know for the exam. Type I errors are those with false positives—that is, you think that evil is present when it is not. You have to be careful with them because if you erroneously raise a red flag and it turns out nothing is wrong, it becomes more difficult to get anyone to listen to you the next time you think you've uncovered something wrong because you've lost credibility.

Type II errors are those with false negatives, where you fail to notice a problem even though it is there—that is, you were looking directly at the evil and didn't recognize it. These errors are generally considered less harmful than Type I errors, though they allow the wrongdoer to get away with the act and maybe even keep doing it.

Type III errors are those in which you come to the right conclusion for all of the wrong reasons. You may conclude that someone broke into your systems because users are having trouble logging in. Someone did indeed break into the system, but you should have noticed it because all of the valuable data is gone.

Risk Management Best Practices

One of the leading ways to address *business continuity* is to do a BIA and implement *best practices*. Best practices are based on what is known in the industry and those methods that have consistently shown superior results over those achieved by other means.

You need only a passing knowledge of business continuity issues for the Security+ exam. If you plan on taking the Project+ exam, also from CompTIA, you will need a more thorough knowledge of these topics.

Undertaking Business Impact Analysis

Business impact analysis (BIA) is the process of evaluating all of the critical systems (important to core business functions) in an organization to define impact and recovery plans. BIA isn't concerned with external threats or vulnerabilities; the analysis focuses on the impact a loss would have on the organization.

Here are the key components of a BIA:

Identifying Critical Functions To identify critical functions, a company must ask itself, "What functions are necessary to continue operations until full service can be restored?" This identification process will help you establish which systems must be returned to operation in order for the business to continue. In performing this identification, you may find that a small or overlooked application in a department may be critical for operations. Many organizations have overlooked seemingly insignificant process steps or systems that have prevented business continuity planning (BCP) from being effective. Every department should be evaluated to ensure that no critical processes are overlooked.

Prioritizing Critical Business Functions When business is continued after an event, operations must be prioritized as *mission-essential* or nonessential functions. If the organization makes resources available to the recovery process, these resources may be limited. Furthermore, in a widespread outage, full operation may not be possible for some time. What would happen, for example, if your data communications services went down? You can usually establish temporary services, but you probably won't be able to restore full network capability. There should be no internal disagreement about the *identification of critical systems*, and you should be clear about which applications or systems have priority based on the resources available. For example, your company may find itself choosing to restore email before it restores its website.

Calculating a Timeframe for Critical Systems Loss How long can the organization survive without a critical function? Some functions in an organization don't require immediate action whereas others do. Which functions must be reestablished and in what timeframe? If your business is entirely dependent on its web presence and is e-commerce oriented, how long can the website stay inoperable? Your organization may need to evaluate and attempt to identify the maximum time that a particular function can be unavailable. When you look at the *impact*, be sure to factor in the following variables:

Life Is anyone in immediate jeopardy because of the failure?

Property Will anything be lost as a result of the malfunction?

Safety Is anyone in harm's way due to the crash?

Finance How much will be lost due to the stoppage?

Reputation How harmful is the breakdown to the trust of the organization?

This component dictates the contingencies that must be established to minimize losses due to exceeding the allowable period.

Estimating the Tangible and Intangible Impact on the Organization Your organization will suffer losses in an outage. These losses will be tangible in nature, such as lost production and lost sales. Intangible losses will also be a factor. For example, will customers lose faith in your service? Knowing the true cost of these impacts in advance will greatly increase the organization's effectiveness in responding to such outages.

A thorough BIA will accomplish several organizational goals:

- The true impact and damage that an outage can cause will be visible.
- Understanding the true loss potential may help you in your fight for a budget.
- Most important, perhaps, the process will document which business processes are being used, the impact they have on the organization, and how to restore them quickly.

The BIA will gain power in the organization as the true costs of an outage become known. People buy insurance not because they intend to have an accident but just in case they have one. A BIA can help identify what insurance is needed for the organization to feel safe.

Identifying Critical Systems and Components

Sometimes your systems are dependent on things that you would not normally consider. Basic utilities such as electricity, water, and natural gas are key aspects of business continuity. In the vast majority of cases, electricity and water are restored—at least on an emergency basis—fairly rapidly. The damage created by blizzards, tornadoes, and other natural disasters is managed and repaired by utility companies and government agencies. Other disasters, such as a major earthquake or hurricane, can overwhelm these agencies, and services may be interrupted for quite a while. When these types of events occur, critical infrastructure may be unavailable for days, weeks, or even months.

Real World Scenario

The Importance of Utilities

When the earthquake of 1989 occurred in San Francisco, California, portions of the city were without electricity, natural gas, and water for several months. Entire buildings were left unoccupied because the infrastructure was badly damaged. This damage prevented many businesses whose information systems departments were located in those buildings from returning to operation for several weeks. Most of the larger organizations were able to shift the processing loads to other companies or divisions.

When you evaluate your business's sustainability, realize that disasters do indeed happen. If possible, build infrastructure that doesn't have a *single point of failure (SPOF)* or connection. After the September 11, 2001, terrorist attack on the World Trade Center (WTC), several ISPs and other companies became nonfunctional because the WTC housed centralized communications systems and computer departments. If you're the administrator for a small company, it is not uncommon for the SPOF to be a router/gateway. The best way to remove an SPOF from your environment is to add redundancy.

Consider the impact of weather on your *contingency plans*. What if you needed to relocate your facility to another region of the country due to a tornado hitting your server room? How would you get personnel there? What personnel would be relocated? How would they be housed and fed during the time of the crisis? You should consider these possibilities in advance. Although the possibility of a crippling disaster is relatively small, you still need to evaluate the risk.

Real World Scenario

Formulating Business Continuity Plans

As a security administrator, you need to think through a way to maintain business continuity should a crisis occur. Imagine that your company is involved in each of the following three scenarios:

Scenario 1

Your company is in the business of monitoring criminal offenders who are under electronic house arrest nationwide. Every offender wears an ankle bracelet that wirelessly communicates with a device in his or her home. The home device communicates to your site in real time over phone lines by calling a toll-free number to report if the offender is in or out of the home, and you alert local authorities immediately if someone isn't in compliance. The number of offenders, and the number of home devices that call your center, is in the tens of thousands. How could business be maintained if the trunk line for the toll-free phone carrier were disrupted in the middle of the night? How could you verify offender compliance if the problem took hours to correct?

Scenario 2

You're the administrator for a small educational company that delivers certification exams locally. The exams are downloaded the night before and delivered throughout the day as students—who have registered over the Internet—arrive. You show up at 8 a.m. on Friday, knowing that there are more than 20 exams to be administered that were downloaded Thursday night. What you find, however, is that someone has broken into the testing room and trashed all of the workstations and monitors. Some of those coming to take the exams are driving from far away. How will you approach the situation?

Scenario 3

You're the database administrator for a large grocery chain. When you leave on Wednesday, there are no problems. When you arrive on Thursday—the day a new sale starts—you learn that the DSL lines are down. They went down before the local stores could download the new sale prices. All scanned goods will ring up at the price they were last week (either sale or regular) and not at current prices. The provider says it's working on the DSL problem but can't estimate how long the repairs will take. How do you approach the problem?

Just like in the real world, there are no right or wrong answers for these scenarios. However, they all represent situations that have happened and for which administrators planned ahead of time.

There are several ways to plan for such scenarios, including implementing redundant technology, fault-tolerant systems, and RAID. A truly redundant system won't use just one of these methods, but rather it will support some aspect of all of them. The following sections address these topics in more detail.

As an administrator, you should always be aware of problems that can occur and have an idea of how you'll approach them. It's impossible to prepare for every emergency, but you can plan for those that could conceivably happen.

Automation/Scripting

The days of relying on someone in the server room to see a problem and push a button to head it off are coming to a close. Thanks to sophisticated monitors and sensors, it is possible to use *automation/scripting* in a wide variety of scenarios to preplan *automated courses of action*. Scenarios in which the automated courses of action can be taken range from *configuration validation* of new equipment added on the network to *continuous monitoring* of server operations.

Frameworks and Templates

Templates can be helpful in the risk assessment process by providing a means to summarize and document results of threat source identification, characterization, vulnerabilities, and impacts. Typical templates include scales for evaluating the threats and deciding the best responses to them.

Master Image

Most newer operating systems allow you to create a model user system as a disk image on a server; the disk image is downloaded and installed when a failure occurs. This method makes it easier for administrators to restore a system than it would be to do it manually.

Nonpersistence

Persistent images are those that stay the same, while nonpersistent are those that are temporary. They can exist only in RAM or be changes that are overwritten on a reboot by a persistent/frozen image.

In a nutshell, a system image (referenced earlier) is a *snapshot* of what exists. Capturing an image of the operating system in its exploited state can be helpful in revisiting the issue after the fact to learn more about it. As an analogy, think of germ samples that are stored in labs after major outbreaks so that scientists can revisit them later and study them further.

Most newer operating systems take snapshots of the configuration at various times and these can be manually created as well (which is highly recommended before updates or major system changes). When something goes awry, you can *revert to known state*—that is, go back to the configuration as it was before the last major change. This ability to *roll back to a known configuration* is helpful on workstations as well as servers.

When all else fails, you can often use *live boot media* to boot a system and begin troubleshooting it. You can usually create bootable flash drives and/or DVDs based on the operating system you are using and the abilities it offers. Bear in mind that for the live boot media to work, the system on which you are working must be configured to boot from that media.

Elasticity

Elasticity is a major feature of cloud computing, meaning the ability to scale up resources as needed. A number of other benefits go along with it: the time to service is a possibility, as is the mean time to implement being quicker inside rather than outside the virtual model, or resource pooling. Other features that make elasticity so valuable include using multitenant models, and the fact that it is scalable not only up but also down, and applications are both available and portable.

Scalability

Speaking of scaling both up and down, *scalability* is always a desired attribute of any system. A virtual datacenter, for example, appears the same as a physical datacenter from an administration standpoint, and it features elasticity, scalability, and so forth. A big benefit of the virtual center is that it can employ a pay-as-you-go model.

Distributive Allocation

Commonly known as load balancing, *distributive allocation* allows for distributing the load (file requests, data routing, and so on) so that no device is overly burdened. This can help with redundancy, availability, and fault tolerance.

High Availability

High availability (HA) refers to the measures, such as redundancy, failover, and mirroring, used to keep services and systems operational during an outage. In short, the goal is to provide all services to all users, where they need them and when they need them. With high availability, the goal is to have key services available 99.999 percent of the time (also known as *five nines availability).*

Planning for Resiliency

Resiliency is the capacity to recover quickly from difficulties. Few things are as difficult in the world of IT as a crash or failure—be it at the server level or the server room level. This chapter is focused on how to compute and manage risk and increase your resiliency when it is economically feasible to do so. For the most part, risk resilience is a buzzword that can be equated with risk management.

Redundancy

Redundancy refers to systems that either are duplicated or *fail over* to other systems in the event of a malfunction. *Failover* refers to the process of reconstructing a system or switching over to other systems when a failure is detected. In the case of a server, the server switches to a redundant server when a fault is detected. This strategy allows service to continue uninterrupted until the primary server can be restored. In the case of a network, this means processing switches to another network path in the event of a network failure in the primary path.

Failover systems can be expensive to implement. In a large corporate network or e-commerce environment, a failover might entail switching all processing to a remote location until your primary facility is operational. The primary site and the remote site would synchronize data to ensure that information is as up to date as possible.

Many operating systems, such as Linux, Windows Server 2012, and Novell Open Enterprise Server, are capable of *clustering* to provide failover capabilities. *Clustering* involves multiple systems connected together cooperatively (which provides *load balancing*) and networked in such a way that if any of the systems fail, the other systems take up the slack and continue to operate. The overall capability of the server cluster may decrease, but the network or service will remain operational.

To appreciate the beauty of clustering, contemplate the fact that this is the technology on which Google is built. Not only does clustering allow you to have redundancy, but it also offers you the ability to scale as demand increases.

Most ISPs and network providers have extensive internal failover capability to provide high availability to clients. Business clients and employees who are unable to access information or services tend to lose confidence. The trade-off for reliability and trustworthiness, of course, is cost: failover systems can become prohibitively expensive. You'll need to study your needs carefully to determine whether your system requires this capability. For example, if your environment requires a high level of availability, your servers should be clustered. This will allow the other servers in the network to take up the load if one of the servers in the cluster fails.

Fault Tolerance

Fault tolerance is the ability of a system to sustain operations in the event of a component failure. Fault-tolerant systems can continue operation even though a critical component, such as a disk drive, has failed. This capability involves over-engineering systems by adding redundant components and subsystems.

Fault tolerance is discussed in more detail in Chapter 7, "Data and Privacy Security Practices," but it appears here as it relates to risk.

Fault tolerance can be built into a server by adding a second power supply, a second CPU, and other key components. Several manufacturers (such as HP, Unisys, and IBM) offer fault-tolerant servers. These servers typically have multiple processors that automatically fail over if a malfunction occurs.

In addition to fault-tolerant servers, you can have fault-tolerant implementations such as Tandem, Stratus, and HP. In these settings, multiple computers are used to provide 100 percent availability of a single server.

There are two key components of fault tolerance that you should never overlook: spare parts and electrical power. Spare parts should always be readily available to repair any system-critical component if it should fail. The redundancy strategy "N+1" means that you have the number of components you need, plus one to plug into any system should it be needed. For example, a small company with five stand-alone servers that are all the same model should have a power supply in a box nearby to install in any one of the servers should there be a failure. (The redundancy strategy 1+1 [or 2N] has one spare part for every component in use.)

Since computer systems cannot operate in the absence of electrical power, it is imperative that fault tolerance be built into your electrical infrastructure as well. At a bare minimum, an *uninterruptible power supply (UPS)*—with surge protection—should accompany every server and workstation. That UPS should be rated for the load it is expected to carry in the event of a power failure (factoring in the computer, monitor, and any other device connected to it) and be checked periodically as part of your preventive maintenance routine to make sure that the battery is operational. You will need to replace the battery every few years to keep the UPS operational.

A UPS will allow you to continue to function in the absence of power for only a short duration. For fault tolerance in situations of longer duration, you will need a *backup generator*. Backup generators run off of gasoline, propane, natural gas, or diesel and generate the electricity needed to provide steady power. Although some backup generators can come on instantly in the event of a power outage, most take a short time to warm up before they can provide consistent power. Therefore, you will find that you still need to implement UPSs within your organization.

Redundant Array of Independent Disks

Redundant Array of Independent Disks (RAID) is a technology that uses multiple disks to provide fault tolerance. There are several designations for RAID levels.

RAID stands for not only *Redundant Array of Independent Disks* but also *Redundant Array of Inexpensive Disks*. Although the latter term has lost its popularity, you might still encounter it in some books.

The most commonly implemented RAID levels are as follows:

RAID Level 0 RAID 0 is *disk striping*. It uses multiple drives and maps them together as a single physical drive. This is done primarily for performance, not for fault tolerance. If any drive in a RAID 0 array fails, the entire logical drive becomes unusable.

RAID Level 1 RAID 1 is *disk mirroring*. Disk mirroring provides 100 percent redundancy because everything is stored on two disks. If one disk fails, another disk continues to operate. The failed disk can be replaced, and the RAID 1 array can be regenerated. This system offers the advantage of 100 percent data redundancy at the expense of doubling the storage requirements. Each drive keeps an exact copy of all information, which reduces the effective storage capability to 50 percent of the overall rated storage. Some implementations of disk mirroring are called *disk duplexing* (*duplexing* is a less commonly used term). The difference between mirroring and duplexing is one more controller card. With mirroring, one controller card writes sequentially to each disk. With duplexing, the same data is written to both disks simultaneously. Disk duplexing has much faster write performance than disk mirroring. Many hardware implementations of RAID 1 are actually duplexing, but they are still generally referred to as mirrors.

The data is intact in a RAID 1 array if either one of the two drives fails. After the failed drive is replaced with a new drive, you remirror the data from the good drive to the new drive to re-create the array.

RAID Level 3 RAID 3 is *disk striping with a parity disk*. RAID 3 arrays implement fault tolerance by using striping (RAID 0) in conjunction with a separate disk that stores parity information. *Parity information* is a value based on the value of the data stored in each disk location. This system ensures that the data can be recovered in the event of a failure. The process of generating parity information uses the arithmetic value of the data binary. This process allows any single disk in the array to fail while the system continues to operate. The failed disk is removed, a new disk is installed, and the new drive is then regenerated using the parity information. RAID 3 is common in older systems, and it's supported by most Unix systems.

RAID Level 5 RAID 5 is *disk striping with parity*, and it is one of the most common forms of RAID in use today. It operates similarly to disk striping, as in RAID 0. The parity

information is spread across all of the disks in the array instead of being limited to a single disk, as in RAID 3. Most implementations require a minimum of three disks and support a maximum of 32.

These four types of RAID drives, or arrays, are illustrated in Figure 1.1.

FIGURE 1.1 The four primary RAID technologies used in systems

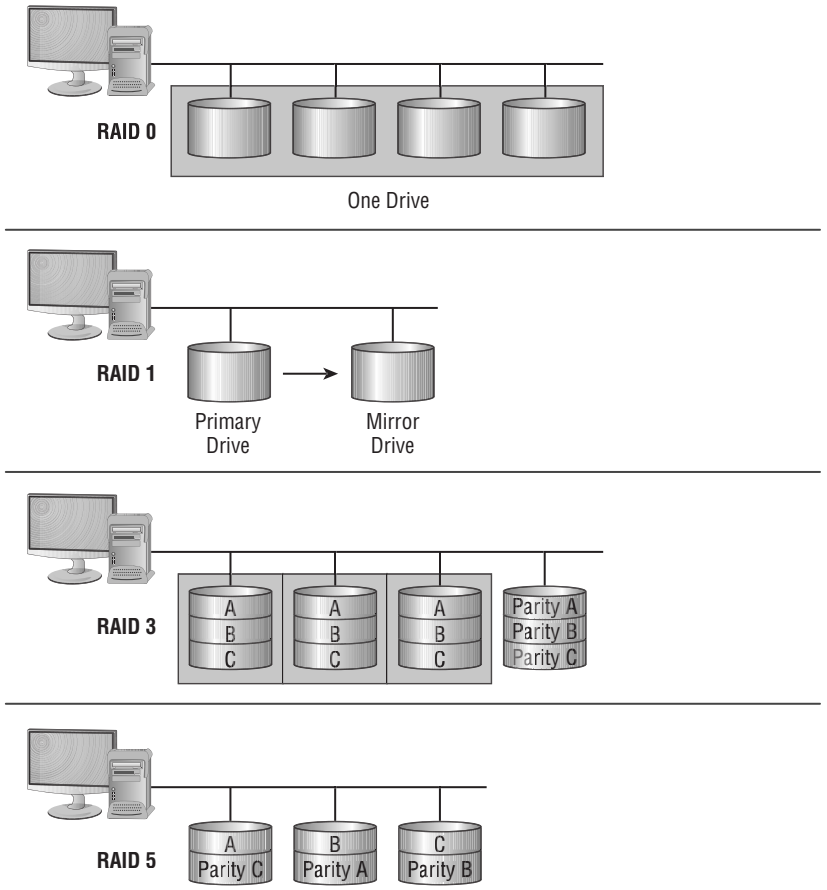

A RAID 5 array can survive the failure of any one drive and still be able to function. It can't, however, survive the failure of multiple drives.

You aren't required to know the current RAID capabilities for the Security+ exam. They are presented here primarily for your edification. They are commonly used in highly reliable systems.

RAID levels 0, 1, 3, and 5 are the ones most commonly implemented in servers today. RAID 5 has largely replaced RAID 3 in newer systems. When two levels are combined for a more potent solution, the numbers simply move into double digits representing the two RAID levels combined. For example, combining RAID 1 with RAID 0 is now called RAID 10 (or RAID 0+1 in older documentation). Combining RAID 1 with RAID 5 is now known as RAID 15, and so on.

RAID levels are implemented either in software on the host computer or in the disk controller hardware. A RAID hardware-device implementation will generally run faster than a software-oriented RAID implementation because the software implementation uses the system CPU and system resources. Hardware RAID devices generally have their own processors, and they appear to the operating system as a single device.

You must do a fair amount of planning before you implement RAID. Within the realm of planning, you must be able to compute the number of disks needed for the desired implementation.

Real World Scenario

How Many Disks Does RAID Need?

As a security administrator, you must determine how many RAID disks you'll need. Compute how many disks will be needed for each of the following scenarios or the amount of storage capacity that results. *(Answers appear after each scenario.)*

Scenario 1

Your company has standardized on 5 TB disks. A new server will go online next month to hold the data files for a new division. The server will be disk-duplexed and needs to be able to store 8 TB of data. How many drives should you order?

Disk duplexing is the same as disk mirroring except that there is also a second controller. Fifty percent of the overall storage capacity must be used for RAID, so you must purchase four 5 TB drives. This will give you excess data capacity of 2 TB.

Scenario 2

Your primary server is currently running four 3 GB disks in a RAID 5 array. Storage space is at a premium, and a purchase order has just been approved for four 5 TB disks. If you still use a RAID 5 array, what is the maximum data storage space this server will be able to host?

The solution that will generate the most data storage capacity is to install all eight drives (the four current ones and the four new ones) into the server. The array must use the same size storage on each drive; thus, all eight drives will appear as if they are 3 TB drives. Under this scenario, 21 TB can be used for data storage, and 3 TB will be used for parity.

Scenario 3

Access speed is of the utmost importance on a web server. You want to purchase some fast 3 TB hard drives and install them in a RAID 0 array. How many drives will you need to purchase to host 900 GB of data?

RAID 0 doesn't perform any fault tolerance and doesn't require any extra disk space. You can obtain 9 TB of data by using three disks.

Change Management

One of the biggest risks an organization faces involves change: either implementing or failing to implement. The discipline of *change management* is focused on how to document and control for change. A subset of project management, change management is focused on controlled implementation and identification of changes. Those changes can be to the actual physical resources themselves (new servers, moving to the cloud, etc.), to individuals (new teams, reorganization, etc.), or even to the organization (merger, acquisition, and so on). The key in every instance is to document everything and focus on the extent and scope of what is affected by every change.

Summary

Risk assessment is the process of evaluating and cataloging the threats, vulnerabilities, and weaknesses that exist in the systems being used. Risk assessment should ensure that all bases are covered.

Security models begin with an understanding of the business issues that an organization is facing. The following business matters must be evaluated:

- Policies
- Standards
- Guidelines

A good policy design includes scope statements, overview statements, accountability expectations, and exceptions. Each of these aspects of a well-crafted policy helps in setting expectations for everyone in a company. For a policy to be effective, it needs the unequivocal support of senior management and decision makers in an organization.

Exam Essentials

Name the three categories of control types. The three types of controls that can be administered are technical, management, and operational.

Know how to calculate risk. Risk can be calculated either qualitatively (subjective) or quantitatively (objective). Quantitative calculations assign dollar amounts, and the basic formula is SLE × ARO = ALE, where SLE is the single loss expectancy, ARO is the annualized rate of occurrence, and ALE is the annual loss expectancy.

Be familiar with the four different approaches to risk. The four risk response strategies are avoidance (don't engage in that activity), transference (think insurance), mitigation (take steps to reduce the risk), and acceptance (be willing to live with the risk).

Know the importance of policies, standards, and guidelines. The process of implementing and maintaining a secure network must first be addressed from a policies, standards, and guidelines perspective. Policies and guidelines set a standard of expectation in an organization. Standards tell people what is expected, and guidelines provide specific advice on how to accomplish a given task or activity.

Understand important elements of key levels of RAID. RAID level 0 does not include any fault tolerance. RAID level 1 can be implemented as mirroring or duplexing; the difference is that the latter includes multiple controllers. RAID level 5 is known as disk striping with parity.

Review Questions

You can find the answers in the Appendix.

1. You're the chief security contact for MTS. One of your primary tasks is to document everything related to security and to create a manual that can be used to manage the company in your absence. Which documents should be referenced in your manual as the ones that identify the methods used to accomplish a given task?
   - **A.** Policies
   - **B.** Standards
   - **C.** Guidelines
   - **D.** BIA

2. Consider the following scenario. The asset value of your company's primary servers is $2 million, and they are housed in a single office building in Anderson, Indiana. Field offices are scattered throughout the United States, but the workstations located at the field offices serve as thin clients and access data from the Anderson servers. Tornados in this part of the country are not uncommon, and it is estimated that one will level the building every 60 years. Which of the following is the SLE for this scenario?
   - **A.** $2 million
   - **B.** $1 million
   - **C.** $500,000
   - **D.** $33,333.33
   - **E.** $16,666.67

3. Refer to the scenario in question 2. Which of the following amounts is the ALE for this scenario?
   - **A.** $2 million
   - **B.** $1 million
   - **C.** $500,000
   - **D.** $33,333.33
   - **E.** $16,666.67

4. Refer to the scenario in question 2. Which of the following is the ARO for this scenario?
   - **A.** 0.0167
   - **B.** 1
   - **C.** 5
   - **D.** 16.7
   - **E.** 60

5. Which of the following strategies involves identifying a risk and making the decision to discontinue engaging in the action?
   - **A.** Risk acceptance
   - **B.** Risk avoidance
   - **C.** Risk mitigation
   - **D.** Risk transference

6. Which of the following policy statements may include an escalation contact in the event that the person dealing with a situation needs to know who to contact?
   - **A.** Scope
   - **B.** Exception
   - **C.** Overview
   - **D.** Accountability

7. Which of the following policies are designed to reduce the risk of fraud and prevent other losses in an organization?
   - **A.** Separation of duties
   - **B.** Acceptable use
   - **C.** Least privilege
   - **D.** Physical access control

8. What is the term used for events that were mistakenly flagged although they weren't truly events about which to be concerned?
   - **A.** Fool's gold
   - **B.** Non-incidents
   - **C.** Error flags
   - **D.** False positives

9. Which of the following is the structured approach that is followed to secure a company's assets?
   - **A.** Audit management
   - **B.** Incident management
   - **C.** Change management
   - **D.** Skill management

10. Which of the following strategies involves sharing some of the risk burden with someone else, such as an insurance company?
    - **A.** Risk acceptance
    - **B.** Risk avoidance
    - **C.** Risk deterrence
    - **D.** Risk mitigation
    - **E.** Risk transference

11. The risk assessment component, in conjunction with the _____, provides the organization with an accurate picture of the situation facing it.

A. RAC

B. ALE

C. BIA

D. RMG

12. Which of the following policy statements should address who is responsible for ensuring that the policy is enforced?

A. Scope

B. Exception

C. Overview

D. Accountability

13. Which of the following strategies is accomplished any time you take steps to reduce risk?

A. Risk acceptance

B. Risk avoidance

C. Risk transference

D. Risk mitigation

14. If you calculate the SLE to be $4,000 and that there will be 10 occurrences a year (ARO), then the ALE is:

A. $400

B. $4,000

C. $40,000

D. $400,000

15. Which of the following policies describes how the employees in an organization can use company systems and resources, both software and hardware?

A. Separation of duties

B. Acceptable use

C. Least privilege

D. Physical access control

16. Separation of duties helps to prevent an individual from embezzling money from a company. To embezzle funds successfully, an individual would need to recruit others to commit an act of _____ (an agreement between two or more parties established for the purpose of committing deception or fraud).

A. Misappropriation

B. Misuse

C. Collusion

D. Fraud

17. Which of the following agreements contains the technical information regarding the technical and security requirements of the interconnection between two or more organizations?

A. BPA

B. MOA

C. ISA

D. MOU

18. If you calculate SLE to be $25,000 and that there will be one occurrence every four years (ARO), then what is the ALE?

A. $6,250

B. $12,500

C. $25,000

D. $100,000

19. Which of the following policies should be used when assigning permissions, giving users only the permissions they need to do their work and no more?

A. Separation of duties

B. Acceptable use

C. Least privilege

D. Physical access control

20. Which of the following strategies necessitates an identified risk that those involved understand the potential cost/damage and agree to live with it?

A. Risk acceptance

B. Risk avoidance

C. Risk transference

D. Risk mitigation