Chapter 1

Today's Information Security Manager

THE CERTIFIED INFORMATION SECURITY MANAGER (CISM) DOMAINS AND SUBTOPICS COVERED IN THIS CHAPTER INCLUDE:

- ✓ Domain 1: Information Security Governance
  - A. Enterprise Governance
    - 1A1. Organizational Culture
    - 1A3. Organizational Structures, Roles and Responsibilities
  - B. Information Security Strategy
    - 1B1. Information Security Strategy Development

THE CERTIFIED INFORMATION SECURITY MANAGER (CISM) SUPPORTING TASKS COVERED IN THIS CHAPTER INCLUDE:

- ✓ 1. Identify internal and external influences to the organization that impact the information security strategy.
- ✓ 2. Establish and/or maintain an information security strategy in alignment with organizational goals and objectives.
- ✓ 7. Gain ongoing commitment from senior leadership and other stakeholders to support the successful implementation of the information security strategy.
- ✓ 8. Define, communicate, and monitor information security responsibilities throughout the organization and lines of authority.

Information security managers are responsible for leading teams of cybersecurity professionals and helping them achieve the goals of the cybersecurity program while aligning those objectives with the needs of the business. This work is crucial to protecting their organizations in today's complex threat landscape. Managers must help their teams protect the confidentiality, integrity, and availability of information and information systems used by their organizations. Fulfilling this responsibility requires a strong understanding of the threat environment facing their organization and a commitment to designing and implementing a set of controls capable of rising to the occasion and answering those threats.

In the first section of this chapter, you will learn about the role that cybersecurity managers play in a modern organization. You will then learn the basic objectives of cybersecurity: confidentiality, integrity, and availability of your operations. In the sections that follow, you will learn about some of the controls that you can put in place to protect your most sensitive data from prying eyes. This chapter sets the stage for the remainder of the book, where you will dive more deeply into many different areas of cybersecurity management.

Information Security Objectives

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

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

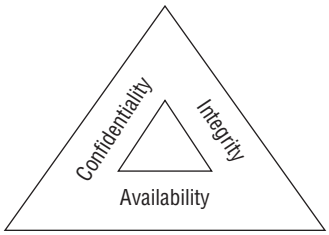

Confidentiality ensures that unauthorized individuals are not able to gain access to sensitive information. Cybersecurity professionals develop and implement security controls, including firewalls, access control lists, and encryption, to prevent unauthorized access to information. Attackers may seek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

Integrity ensures that there are no unauthorized modifications to information or systems, either intentionally or unintentionally. Security professionals use integrity controls, such as hashing and integrity monitoring solutions, to enforce this requirement. Integrity threats may come from attackers actively seeking the alteration of information without authorization, or they may result from human error, mechanical failure, or environmental conditions, such as a power spike corrupting information.

Availability ensures that information and systems are ready to meet the needs of legitimate users at the time those users request them. Security professionals use availability controls, such as fault tolerance, clustering, and backups, to ensure that legitimate users gain access as needed. Similar to integrity threats, availability threats may come from attackers actively seeking the disruption of access, or they may come from human error, mechanical failure, or environmental conditions, such as a fire destroying a data center that contains valuable information or services.

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

Role of the Information Security Manager

Information security managers are responsible for safeguarding the confidentiality, integrity, and availability of the information and systems used by their organization. But they must achieve these goals within the context of the organization's day-to-day activities and strategic objectives. The information security manager must wear the two hats shown in Figure 1.2: that of a cybersecurity subject matter expert and that of a business leader engaged with the organization's mission.

FIGURE 1.2 Information security managers must be both security experts and business leaders.

This "dual-hattedness" is perhaps the most significant defining characteristic of what makes an information security leader different from an information security professional. Information security professionals can narrow much of their focus to cybersecurity matters. Leaders, on the other hand, must maintain that organizational focus at the same time and use their expertise to help guide the organization in making decisions that are both sound from a business perspective and reasonable from a risk management perspective.

Depending on the size of an organization, information security management and leadership may be a role shared by several (or many!) different people, a consolidated role held by a single person, or even a partial role filled by someone who also bears other responsibilities within the organization. There is no one-size-fits-all answer to sizing the information security function for an organization—the selection is highly dependent on the nature of the organization's security requirements, the complexity of their operating environment, and the team they have in place.

Chief Information Security Officer

The most senior information security leader within an organization often bears the title of *chief information security officer (CISO)*. The CISO is a senior business executive who is responsible for overseeing all information security efforts within the organization. The CISO title is commonly accepted as the standard for an organization's information security leader, although some organizations may use different titles, including these:

- Vice president for information security (or assistant/associate vice president)
- Director of information security
- Information security manager

Many people believe that the use of these alternative titles indicates diminished status in the organization and a lack of prioritization for cybersecurity. In many cases, there is some truth behind this perception. In some cases, the use of the term *officer* may also imply that the individual bearing the title is an officer of the corporation or nonprofit organization. This has specific legal consequences that affect the CISO's responsibility and personal liability. However, it is important to note that just because someone has the title of CISO does not automatically make them an officer of the organization. Election or appointment as an officer is a formal process that requires the consent of the governing board.

The choice of a title also varies widely based on industry practices and organizational culture. For our purposes, we will continue to refer to the senior-most information security leader as the CISO throughout this book.

Lines of Authority

The lines of authority for the CISO also convey the role that cybersecurity plays in the organization, both the number and the functions of people reporting to the CISO, and the person to whom the CISO reports. It is quite common for the CISO to report to the chief information officer (CIO), who leads the IT function. This CISO/CIO reporting relationship

clearly places responsibility for information security issues within the IT organization. In other cases, the CISO may report to other executives, such as:

- Chief executive officer (CEO)
- Chief risk officer (CRO)
- Chief security officer (CSO) (this role includes oversight of information security, physical security, and other security concerns)
- Chief operating officer (COO)
- Chief audit executive

The nature of this reporting relationship signals the importance that the organization places on the cybersecurity function as well as the perceived role of cybersecurity within the organization. For example, placing the information security function underneath a chief risk officer or chief security officer signals that the organization views information security risks within the context of a broader enterprise risk management or security program. As with titling, there is no "correct" placement of the CISO within the organizational structure, but organizations should be cognizant of the message they send to the security team and other employees based on their selection.

Although there are strong arguments for placing information security in several different parts of the organization, one general principle that should almost always be observed is that information security should not be buried underneath another function. This is particularly true when doing so may create a conflict of interest. For example, an organization might decide to place the information security function under a director of technology infrastructure who reports to the CIO. This approach is problematic for several reasons:

- It indicates that information security is not as important to the organization as other technology functions.
- It creates a potential conflict of interest when the information security team disagrees with an approach endorsed by the director to whom they report or when the security team is expected to report unflattering results of audits or tests performed on assets owned by that individual.
- It creates difficulties when the cybersecurity team has a conflict with another technology team that resides in a different part of the IT organization.

Organizing the Security Team

The CISO bears ultimate responsibility for protecting the confidentiality, integrity, and availability of the information and systems used by the organization. The specific controls and techniques used to achieve those goals will vary greatly, depending on the nature of the organization and its security requirements.

In almost every case, a team of information security professionals supports the CISO in their work, providing subject matter expertise and operational talents to achieve the organization's security objectives. In larger organizations, the CISO may leverage a management structure similar to the one shown in Figure 1.3, where a director who reports to the CISO

leads each major cybersecurity function. The specific functions shown in Figure 1.3 are for illustrative purposes only and will vary from organization to organization.

FIGURE 1.3 Typical cybersecurity organizational structure

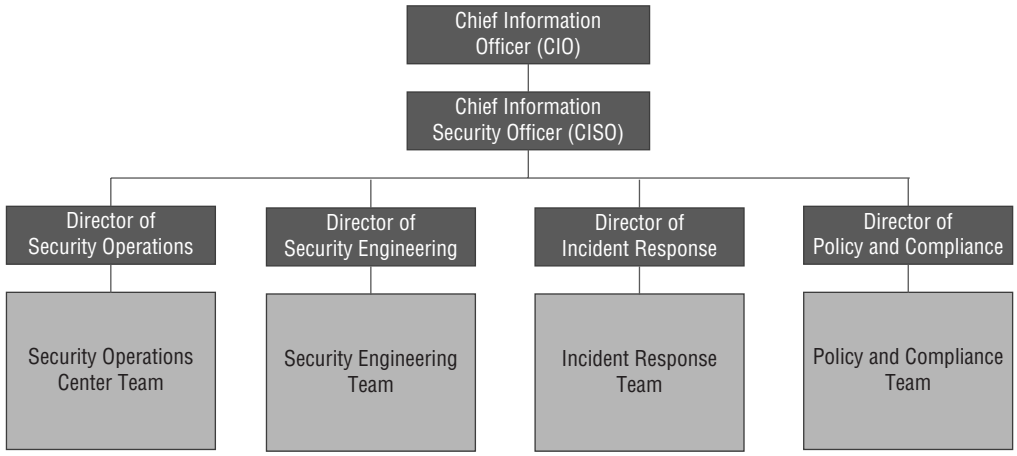

In larger organizations, these directors may each be supported by a series of managers, each of whom has individual contributors as direct reports. In midsized organizations, individual contributors may report to the directors. The need for additional layers of management is directly dependent on the number of people in the organization and should be optimized to reflect each manager's span of control. The *span of control* represents the number of individuals who directly report to a position. Different organizations have different philosophies on span of control, but it is commonly thought that managers with less than five direct reports likely have too small of a span of control and could take on additional responsibilities, whereas managers with more than 10 direct reports may have difficulty effectively managing a very large team.

Cybersecurity vs. Information Security

You may have noticed that we use the terms *cybersecurity* and *information security* almost interchangeably throughout this book. This is a deliberate choice, because the terms are commonly used this way in practice. However, we do want to point out that they do have two different meanings.

As you've already read, the goal of information security is to protect the confidentiality, integrity, and availability of an organization's information and information assets. The goal

of cybersecurity is to protect the confidentiality, integrity, and availability of an organization's digital resources. These terms are closely related, but there are subtle differences.

Information security, properly defined, is responsible for the security of *all* information, whether in digital or analog form. An information security program would be responsible not only for electronic information systems but also for the protection of paper records and other nondigital assets. Cybersecurity, properly defined, is responsible for the security of all digital assets and may be thought of as a subset of the field of information security.

Although these definitional differences exist, the reality is that cybersecurity teams are almost always responsible for information security more broadly and most practitioners use these terms interchangeably. Therefore, for the sake of variety, we will do so as well in this book.

Roles and Responsibilities

Responsibility for different information security functions may be spread among a team and across the organization. For example, consider an organization's response to a cybersecurity incident. The organization may decide that the CISO has overall accountability for the incident response effort. However, the CISO does not do this on their own. They are supported by a variety of stakeholders who play different roles. The incident response team leader and members report to the CISO and carry out the actual response. Legal counsel provides valuable input on compliance issues and responsibilities. The CEO may need to be kept informed of incident progress. Tracking all of these stakeholders is crucial to ensuring that items don't slip through the cracks.

The RACI matrix is a common management tool used to specify how roles and responsibilities are shared throughout an organization. The matrix shows various security responsibilities and roles and then includes one of four letters indicating the level of involvement each role has in that responsibility. The options for filling in the RACI matrix are as follows:

- *Responsible (R)* roles are those who actually carry out the work involved. There must be at least one role assigned as responsible for each responsibility, although there may be more than one.
- *Accountable (A)* roles bear ultimate and final responsibility for achieving the objective. Consider this the "buck stops here" role for the responsibility. Each responsibility in the matrix must have one, and only one, accountable role.
- *Consulted (C)* roles are those who provide input that affects the responsibility because of their subject matter expertise.
- *Informed (I)* roles are those who are provided with regular updates on the status of the effort. They may need this information to complete their work, oversee the organization, or perform other tasks, but the key characteristic is that, unlike consulted roles, informed roles receive updates but do not provide input.

Figure 1.4 shows an abbreviated example of a RACI matrix for a few security roles in an organization.

FIGURE 1.4 RACI matrix for information security

	Incident Response	Privacy Compliance	Security Leadership
CEO	I		I
CIO	I	A	C
CISO	A	R	A
IR Leader	R	C	R
IR Team	R	C	
Compliance Leader	C	R	R
Compliance Team		R	
SOC Leader	C	I	R
SOC Team	C		
Legal Counsel	C	R	
Public Relations	C		

Information Security Risks

Security incidents occur when an organization experiences an adverse impact to the confidentiality, integrity, and/or availability of information or information systems. These incidents may occur as the result of malicious activity (such as an attacker targeting the organization and stealing sensitive information); accidental activity (such as an employee leaving an unencrypted laptop in the back of a rideshare); or natural activity (such as an earthquake destroying a data center).

Security professionals are responsible for understanding these risks and implementing controls designed to manage those risks to an acceptable level. To do so, they must first understand the effects that an incident might have on the organization and the impact it might have on an ongoing basis.

The DAD Triad

Earlier in this chapter, we introduced the CIA Triad, used to describe the three main goals of cybersecurity: confidentiality, integrity, and availability. Figure 1.5 shows a related model: the *DAD Triad*. This model explains the three important threats to cybersecurity efforts: *disclosure*, *alteration*, and *denial*. Each of these three threats maps directly to one of the main goals of cybersecurity:

- *Disclosure* is the exposure of sensitive information to unauthorized individuals, otherwise known as *data loss*. Disclosure is a violation of the principle of confidentiality. Attackers who gain access to sensitive information and remove it from the organization are said to be performing *data exfiltration*. Disclosure may also occur accidentally, such as when an administrator misconfigures access controls or an employee loses a device.
- *Alteration* is the unauthorized modification of information and is a violation of the principle of integrity. Attackers may seek to modify records contained in a system for financial gain, such as adding fraudulent transactions to a financial account. Alteration may occur as the result of natural activity, such as a power surge causing a "bit flip"

that modifies stored data. Accidental alteration is also a possibility, if users unintentionally modify information stored in a critical system as the result of a typo or other unintended activity.

- *Denial* is the disruption of an authorized user's legitimate access to information. Denial events violate the principle of availability. This availability loss may be intentional, such as when an attacker launches a distributed denial-of-service (DDoS) attack against a website. Denial may also occur as the result of accidental activity, such as the failure of a critical server, or as the result of natural activity, such as a natural disaster impacting a communications circuit.

FIGURE 1.5 The three key threats to cybersecurity programs are disclosure, alteration, and denial.

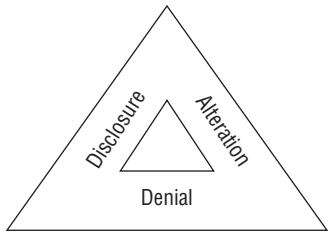

The CIA and DAD triads are very useful tools for cybersecurity planning and risk analysis. Whenever you find yourself tasked with a broad goal of assessing the security controls used to protect an asset or the threats to an organization, you can turn to the CIA and DAD triads. For example, if you're asked to assess the threats to your organization's website, you may apply the DAD Triad in your analysis:

- Does the website contain sensitive information that would damage the organization if disclosed to unauthorized individuals?
- If an attacker were able to modify information contained on the website, would this unauthorized alteration cause financial, reputational, or operational damage to the organization?
- Does the website perform mission-critical activities that could damage the business significantly if an attacker were able to disrupt the site?

That's just one example of using the DAD Triad to inform a risk assessment. You can use the CIA and DAD models in almost any situation to serve as a helpful starting point for a more detailed risk analysis.

Incident Impact

The impacts of a security incident may be wide-ranging, depending on the nature of the incident and the type of organization affected. We can categorize the potential impact of a security incident using the same categories that businesses generally use to describe any type of risk: financial, reputational, strategic, operational, and compliance.

Let's explore each of these risk categories a little further.

Financial Risk

Financial risk is, as the name implies, the risk of monetary damage to the organization as the result of a data breach, service disruption, or other security incident. This may be very direct financial damage, such as the costs of rebuilding a data center after it is physically destroyed or the costs of contracting experts for incident response and forensic analysis services.

Financial risk may also be indirect and come as a second-order consequence of the breach. For example, if an employee loses a laptop containing plans for a new product, the organization suffers direct financial damages of a few thousand dollars from the loss of the physical laptop. However, the indirect financial damage may be more severe—competitors may get ahold of those product plans and beat the organization to market, resulting in potentially significant revenue loss.

Reputational Risk

Reputational risk occurs when the negative publicity surrounding a security breach causes the loss of goodwill among customers, employees, suppliers, and other stakeholders. It is often difficult to quantify reputational damage, since these stakeholders may not directly say that they will reduce or eliminate their volume of business with the organization as a result of the security breach. However, the breach may still have an impact on their future decisions about doing business with the organization.

Identity Theft

When a security breach strikes an organization, the effects of that breach often extend beyond the walls of the breached organization, affecting customers, employees, and other individual stakeholders. The most common impact on these groups is the risk of identity theft posed by the exposure of personally identifiable information (PII) to unscrupulous individuals.

Organizations should take special care to identify, inventory, and protect PII elements, especially those that are prone to use in identity theft crimes. These include Social Security numbers, bank account and credit card information, driver's license numbers, passport data, and similar sensitive identifiers.

Strategic Risk

Strategic risk is the risk that an organization will become less effective in meeting its major goals and objectives as a result of the breach. Consider again the example of an employee losing a laptop that contains new product development plans. In addition to the financial impact discussed earlier, this incident may pose strategic risk to the organization in two different ways. First, if the organization does not have another copy of those plans, they may

be unable to bring the new product to market or may suffer significant product development delays. Second, if competitors gain hold of those plans, they may be able to bring competing products to market more quickly or even beat the organization to market, gaining first-mover advantage. Both of these effects demonstrate strategic risk to the organization's ability to carry out its business plans.

Operational Risk

Operational risk is risk to the organization's ability to carry out its day-to-day functions. Operational risks may slow down business processes, delay delivery of customer orders, or require the implementation of time-consuming manual workarounds to normally automated practices.

Operational risk and strategic risk are closely related, so it might be difficult to distinguish between them. Think about the difference in terms of the nature and degree of the impact on the organization. If a risk threatens the very existence of an organization or the ability of the organization to execute its business plans, that is a strategic risk that seriously jeopardizes the organization's ongoing viability. On the other hand, if the risk only causes inefficiency and delay within the organization, it fits better into the operational risk category.

Compliance Risk

Compliance risk occurs when a security breach causes an organization to run afoul of legal or regulatory requirements. For example, the Health Insurance Portability and Accountability Act (HIPAA) requires that health-care providers and other covered entities protect the confidentiality, integrity, and availability of protected health information (PHI). If an organization loses patient medical records, they run afoul of HIPAA requirements and are subject to sanctions and fines from the U.S. Department of Health and Human Services. That's an example of compliance risk.

Organizations face many different types of compliance risk in today's regulatory landscape. The nature of those risks depends on the jurisdictions where the organization operates, the industry that the organization functions within, and the types of data that the organization handles. We discuss these compliance risks in more detail in Chapter 2, "Information Security Governance and Compliance."

Risks Often Cross Categories

Don't feel like you need to shoehorn every risk into one and only one of these categories. In most cases, a risk will cross multiple risk categories. For example, if an organization suffers a data breach that exposes customer PII to unknown individuals, the organization will likely suffer reputational damage due to negative media coverage. However, the organization may also suffer financial damage. Some of this financial damage may come in the form of lost business due to the reputational damage. Other financial damage may come

(continues)

(continued)

as a consequence of compliance risk if regulators impose fines on the organization. Still more financial damage may occur as a direct result of the breach, such as the costs associated with providing customers with identity protection services and notifying them about the breach.

Building an Information Security Strategy

Perhaps the most important responsibility of an information security leader is the creation, implementation, and maintenance of an *information security strategy* for the organization. This strategy begins with an assessment of the current state of the organization and a comparison to the desired state of security based on the organization's control objectives. It then outlines a plan for working from that current state to achieve the desired state through clearly articulated goals.

Before developing an information security strategy, information security leaders should gather information about the current and desired states of the organization. They do this through a series of analyses, including threat research, SWOT analysis, and gap analyses.

Threat Research

Developing a cybersecurity strategy requires a strong understanding of the threat environment facing cybersecurity professionals. A strategy is only effective if it combats the threats that pose the greatest risk to the organization. These threats may be described using two important factors:

- *Threat actors* are the individuals or groups seeking to undermine the security of an organization.
- *Threat vectors* are the tactics, tools, and techniques used by threat actors to achieve their objectives.

Cybersecurity threat actors differ significantly in their skills, capabilities, resources, and motivation. Protecting your organization's information and systems requires a solid understanding of the nature of these different threats so that you may develop a set of security controls that comprehensively protect your organization against them.

We dedicate an entire chapter of this book to understanding the cybersecurity threat landscape and conducting threat research. You will learn more about these topics in Chapter 4, "Cybersecurity Threat."

SWOT Analysis

SWOT analysis is a technique commonly used by organizations to assess their current state and develop their forward-looking strategy. SWOT is an acronym describing the four major elements of the analysis:

- *Strengths* are internal characteristics of the organization that provide it with an advantage toward achieving its goals/mission. For example, a cybersecurity team might consider its cybersecurity awareness program as a strength if it is particularly effective.
- *Weaknesses* are internal characteristics of the organization that place it at a disadvantage toward achieving its goals/mission. For example, a cybersecurity team might identify the lack of application security skills as a weakness.
- *Opportunities* are external factors that the organization might exploit to better achieve its goals/mission. For example, a cybersecurity team might consider the use of managed service providers as an opportunity to relieve the burden on the team and focus their work on value-added activities.
- *Threats* are external factors that might jeopardize the organization's ability to achieve its goals/mission. For example, a new privacy law passed by a jurisdiction within which the company operates might pose a threat to the organization.

The SWOT analysis may be conducted at any level of the organization. Senior leaders may conduct a SWOT analysis that analyzes the business overall. The CISO may conduct a SWOT analysis for the broad information security function, whereas the director of the incident response team may conduct a SWOT analysis for that specific function.

Organizations typically develop a SWOT analysis through a collaborative process that seeks inputs from all levels of the team, from individual contributors to senior management. The exercise of creating a SWOT analysis helps the organization think critically about its current position and how it will be affected by both internal and external forces moving forward.

A SWOT analysis may be quite detailed, but teams usually document the result of their work in a generalized chart similar to the one shown in Figure 1.6. This matrix organizes positive factors (strengths and opportunities) on the left side and negative factors (weaknesses and threats) on the right side. Similarly, internal factors (strengths and weaknesses) appear on the top, and external factors (opportunities and threats) appear on the bottom.

Gap Analysis

After identifying the risks that they face, organizations define their security requirements by writing a series of *control objectives* that describe how they plan to manage those risks. These control objectives are described from a strategic perspective in a general manner and provide a basis for the evaluation of the organization's current information security program against its desired state.

FIGURE 1.6 Cybersecurity SWOT analysis example

	Positive	Negative
Internal	**Strengths** 1. Experienced team 2. Strong technology infrastructure 3. Ability to innovate	**Weaknesses** 1. Incident response skills 2. Disorganized vendor management process 3. Lack of consolidated logging
External	**Opportunities** 1. Managed service provider offerings 2. Vendor-provided training	**Threats** 1. New state-level compliance requirements 2. Security failures at service providers

Control Objectives

In many cases, organizations draw these control objectives from industry standard frameworks, such as the Control Objectives for Information Technology (COBIT). Developed by ISACA, COBIT provides broad objective statements that apply to any IT organization. For example, the COBIT control objective for managed security states:

> Keep the impact and occurrence of information security incidents within the enterprise's risk appetite levels.

You'll learn more about the COBIT framework and other approaches to developing control objectives in Chapter 2.

With those control objectives in hand, cybersecurity managers can conduct an assessment of the current state of their controls and determine the degree to which they are achieving their control objectives. This process, known as a *gap analysis*, identifies areas of deficiency and opportunities for improvement that, if prioritized for remediation, may become the basis for goals in the organization's information security strategy.

Maturity Models

In addition to performing an objectives-based gap analysis, organizations may use *maturity models* to assess the state of their IT organization against industry best practices. ISACA

offers the *Capability Maturity Model Integration (CMMI)* as a method to assess maturity of an organization. The use of these models is particularly common in software development efforts and in U.S. government contracting work, but the model may also be applied to security and other processes.

FIGURE 1.7 CMMI levels

Under CMMI, an organization assesses each process as being at one of the five levels shown in Figure 1.7:

- At *Level 1: Initial*, the organization has unpredictable processes that are poorly controlled. This level is characterized by reactive management and a "firefighting" approach.
- When an organization achieves *Level 2: Managed*, it begins to implement organized processes on a per-project basis but is still operating in reactive mode.
- Moving on to *Level 3: Defined*, the organization has standard processes that are used organization wide and are adapted for use within each project. This level marks a shift from reactive to proactive management.
- *Level 4: Quantitatively Managed* organizations build measurement and controls on top of their processes to allow them to quickly identify and remediate deficiencies and address control gaps before issues arise.
- At the top tier of the CMMI, *Level 5: Optimizing* organizations use a continuous process improvement approach to adjust and fine-tune the way that they work to achieve peak efficiency and effectiveness.

Creating SMART Goals

Drafting the goals for an information security strategy can be a daunting task. Information security managers need to create goals that motivate the organization to succeed, move the organization toward the desired security state, and are reasonable to achieve. Strategies should not be vague but rather clearly articulated. For example, a goal of "Improve our vulnerability management program" is difficult to specifically envision. What types of improvement are necessary? When will they be completed? How will the organization know whether it has achieved the goal?

Organizations developing clearly articulated goals often use the SMART framework to describe the characteristics each goal should possess. The five characteristics of a SMART goal are as follows:

- The goal is *specific*. It describes clearly what the organization intends to achieve.
- The goal is *measurable*. It includes clear criteria by which the organization can measure success.
- The goal is *achievable*. The organization can realistically achieve the goal within the specified time period.
- The goal is *relevant*. If achieved, the goal will advance the organization's strategic objectives.
- The goal is *time-bound*. It includes a specific deadline for achievement.

Let's revisit the vague goal of "Improve our vulnerability management program." We can transform this goal into a SMART goal by recasting it as "Implement daily network vulnerability scanning for all production systems by the end of this year." This goal is much more specific—we are going to conduct network vulnerability scans on a daily basis. It is also measurable. We will know that we have achieved the goal once we are scanning 100 percent of production systems on a daily basis. Whether the goal is achievable within the next year or relevant to the organization's strategy is a judgment that must be made by management.

Don't confuse "achievable" with "easy." Leaders should push their teams to achieve difficult goals by assigning stretch goals that force teams out of their comfort zone to achieve higher levels of productivity. At the same time, teams should not be set up for failure. Balancing the appropriate level of stress to place on the organization is an important job of the information security manager.

Alignment with Business Strategy

Information security functions exist for only one purpose: to serve the business. Certainly, security teams are focused on protecting the confidentiality, integrity, and availability of that business's information and systems, but information security managers must remain constantly aware that they do so in service of the organization achieving its business goals and objectives.

This is often one of the most important challenges facing leaders of cybersecurity teams. It's easy for technical subject matter experts to get lost in the weeds of their work and come to think of cybersecurity as an end in and of itself, but cybersecurity is only effective when it facilitates the achievement of organizational goals and objectives. Information security efforts must align with the business's goals, objectives, functions, processes, and practices.

Leadership Support

As a supporting function, information security initiatives do not generate revenue. Security functions are a cost center from a financial perspective. Every dollar spent on cybersecurity issues is a dollar that cannot be invested elsewhere in the business or returned to shareholders as profit. Therefore, senior business leaders and other stakeholders are often wary of investments in cybersecurity, and achieving their support is crucial to the success of the program.

One of the most important responsibilities of the information security manager is to gain ongoing commitment from senior leadership and other important stakeholders for investments of time and money in the security program. This requires helping leaders understand how information security efforts support the organization's goals and objectives. It also requires developing business cases for cybersecurity initiatives that demonstrate the impact of those initiatives on the business and include clear criteria for determining their successful implementation.

You'll learn more about building business cases for cybersecurity in Chapter 2.

Internal and External Influences

As you develop a security strategy for your organization, your core task is to achieve your security objectives in a manner that aligns with your organization's business strategy. That said, it would be naïve to believe that this happens in a bubble. Your organization's approach to security is influenced by a number of internal and external factors:

- The broader *business environment* within which your organization operates. The demands of customers and pressures placed on the organization by competitors will influence the level of commitment to and investment in cybersecurity, both positively and negatively.
- Your organization's *risk tolerance*. This is the degree of risk that you are willing to undertake as you seek to achieve your business objectives. You will learn more about risk tolerance and risk management in Chapter 3, "Information Risk Management."
- The *regulatory environment* within which your organization operates. This may include federal and state laws and industry codes of practice. You will learn more about regulatory requirements that may apply to your organization in Chapter 2.
- Changes in the *threat landscape*. As adversaries adapt their tactics and techniques, your cybersecurity strategy must evolve to combat those changes. You will learn more about the cybersecurity threat landscape in Chapter 4.

- *Emerging technologies* in use in your field. If competitive pressures or innovation strategies guide your organization toward the use of emerging technologies, those technologies will challenge the status quo in your security program.
- *Social media* spreads news at faster rates than ever before. Even if your organization does not directly discuss cybersecurity issues on social media, rest assured that your customers and other stakeholders will.
- *Third-party considerations* also play a role. Although the media, industry groups, vendors, and other third parties may not have regulatory authority, they may also bring pressure to bear on your organization, forcing changes in cybersecurity strategies.

Cybersecurity Responsibilities

You've heard the old adage: security is everybody's responsibility. There's wisdom in that old saying—cybersecurity professionals aren't the only ones who must protect the organization's information and information systems. As you build out your information security strategy, be sure to clearly document the roles of major contributors. These include three critical roles in data governance: data owners, data stewards, and data custodians.

Data owners are the senior-level officials who bear overall responsibility for particular datasets. The data owner sets policies and guidelines for data use and data security and has the authority to make final decisions regarding a dataset. Data owners are usually the business leaders who have responsibility for the mission area most closely related to the dataset. For example, an organization's vice president for human resources might be the data owner for employment information.

Practically speaking, most individuals who are senior enough to hold the position of data owner do not have the time available to get involved in the daily decisions of data governance. They usually delegate that responsibility to a *data steward*. The data steward handles the implementation of the high-level policies set by the data owner. For example, a data steward might make day-to-day decisions about who may access a dataset. In the case of the employee dataset, if the data owner is the vice president for human resources, that vice president might delegate data stewardship responsibility to a director for HR information services. In most cases, there is a reporting relationship between the data owner and the data steward.

*Data custodian*s are the individuals who actually store and process the information in question. IT staff often find themselves in the position of data custodians because of their roles as system owners and administrators. Technologists are rarely data owners or data stewards, but they are usually data custodians for almost all of the data in the organization due to the nature of their jobs. Data stewards ensure that appropriate data protections are in place, including encryption, backups, access controls, and other mechanisms that meet the requirements set forth by data owners and stewards.

Data processors are third-party organizations that handle data on behalf of a data owner. For example, if the IT team at an organization stores data in a cloud service, that cloud service provider is a data processor.

Consider an example that helps tie these terms together. If your bank collects financial information from you to process loan applications and an IT administrator at the bank uses a cloud service to store those records, we have several roles at play. You are the data subject, because the records are about you. The bank, as an organization, is responsible for that data and likely designates a senior officer, such as the vice president of loans, as the data owner. The IT administrator who handles the records is a data custodian, and the cloud service they use is a data processor.

Individual users also bear responsibility for protecting the security of information and systems that they use and access. Cybersecurity responsibility training should be provided to all end users but should also have a particular focus on two categories of user:

- *High-risk users* who are the likely targets of cyberattacks. This may be because they are high-profile individuals likely to attract attention or because they engage in activities that place them at higher risk, such as frequently traveling to high-risk destinations.
- *Privileged users* who would pose a higher-than-average risk if their accounts were compromised. This includes technologists with administrative access to systems, finance professionals with the ability to initiate funds transfers, executives with access to sensitive information, and other similar highly privileged groups.

Strong cybersecurity programs clearly define the responsibilities of each of these groups, communicate to them regularly, and monitor their progress toward achieving security objectives.

Communication

Cybersecurity strategies are only effective if they are clearly communicated to stakeholders throughout the organization. You will need to use messaging and methods for this communication that fit within your organization's normal business processes. You will want to consider the normal culture within your organization and traditional channels of communication and take advantage of those as much as possible. For example, if employees are used to receiving important messages at their weekly team meetings, try to inject security messaging into those meetings. If they prefer brief informal communication by Slack, communicate that way. Bring the cybersecurity message to users where they already are and they'll be much more receptive to that messaging.

At the same time, cybersecurity messages must be concise. Highlight the essential aspects of information security in a way that translates to actionable advice for the audience. They don't need to hear all of the details behind cybersecurity strategies. They just need to understand that their assistance is important and what, exactly, is being asked of them.

Action Plans

As you put a cybersecurity strategy in place, you'll want to develop an action plan that outlines both the short-term and long-term steps that you will take to move your organization from its current state to its desired state. As with communication efforts, you'll be more

successful if you align those plans with the normal project planning methods used by your organization.

Figure 1.8 shows an example of a one-page plan that covers the four-year rollout of an organization's information security strategy.

FIGURE 1.8 Communicating the security strategy

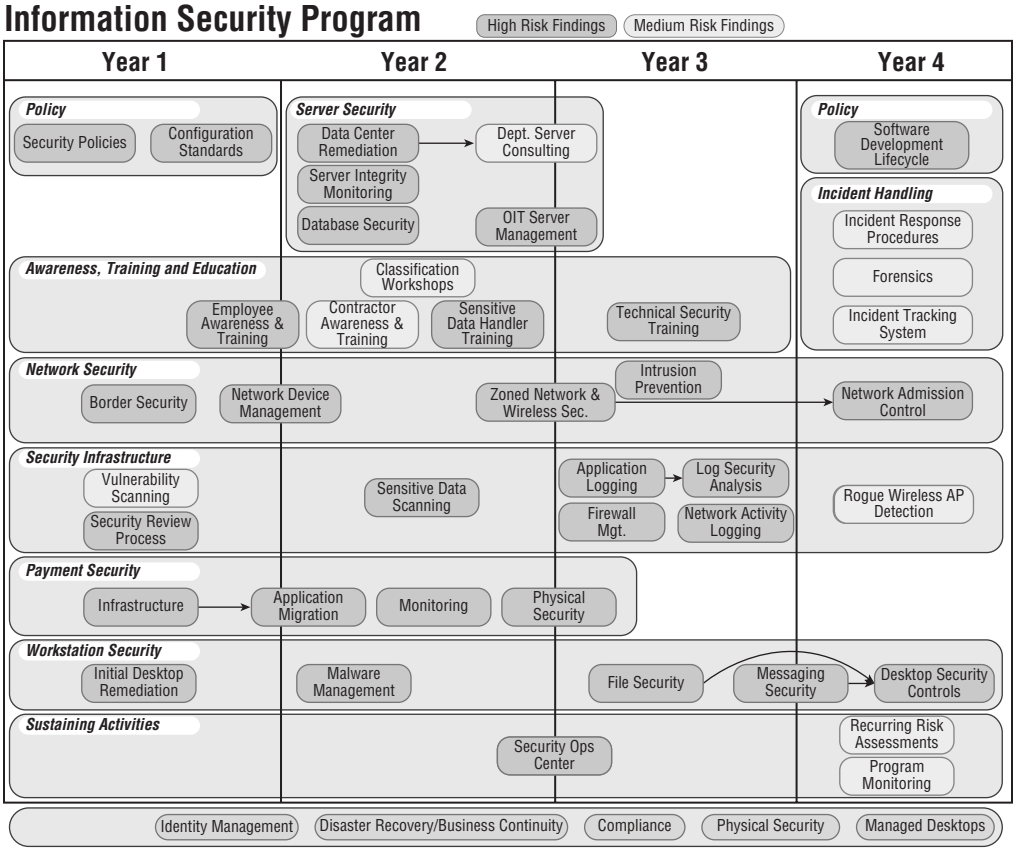

Implementing Security Controls

As an organization analyzes its risk environment, technical and business leaders determine the level of protection required to preserve the confidentiality, integrity, and availability of their information and systems. They express these requirements by writing the *control*

objectives that the organization wishes to achieve. These control objectives are statements of a desired security state, but they do not, by themselves, actually carry out security activities. *Security controls* are specific measures that fulfill the security objectives of an organization.

Security Control Categories

Security controls are categorized based on their mechanism of action: the way that they achieve their objectives. There are three different categories of security control:

- *Technical controls* enforce confidentiality, integrity, and availability in the digital space. Examples of technical security controls include firewall rules, access control lists, intrusion prevention systems, and encryption.
- *Operational controls* include the processes that we put in place to manage technology in a secure manner. These include user access reviews, log monitoring, and vulnerability management.
- *Managerial controls* are procedural mechanisms that focus on the mechanics of the risk management process. Examples of administrative controls include periodic risk assessments, security planning exercises, and the incorporation of security into the organization's change management, service acquisition, and project management practices.

If you're not familiar with some of the controls provided as examples in this chapter, don't worry about it! We'll discuss them all in detail later in the book.

Organizations should select a set of security controls that meets their control objectives based on the criteria and parameters that they either select for their environment or have imposed on them by outside regulators. For example, an organization that handles sensitive information might decide that confidentiality concerns surrounding that information require the highest level of control. At the same time, they might conclude that the availability of their website is not of critical importance. Given these considerations, they would dedicate significant resources to the confidentiality of sensitive information while perhaps investing little, if any, time and money protecting their website against a denial-of-service attack.

Many control objectives require a combination of technical, operational, and management controls. For example, an organization might have the control objective of preventing unauthorized access to a data center. They might achieve this goal by implementing biometric access control (technical control), performing regular reviews of authorized access (operational control), and conducting routine risk assessments (managerial control).

Security Control Types

We can also divide security controls into types, based on their desired effect. The types of security control include the following:

- *Preventive controls* intend to stop a security issue before it occurs. Firewalls and encryption are examples of preventive controls.

- *Detective controls* identify security events that have already occurred. Intrusion detection systems are detective controls.
- *Corrective controls* remediate security issues that have already occurred. Restoring backups after a ransomware attack is an example of a corrective control.
- *Deterrent controls* seek to discourage an attacker from attempting to violate security policies. Vicious guard dogs and barbed wire fences are examples of deterrent controls.
- *Physical controls* are security controls that impact the physical world. Examples of physical security controls include fences, perimeter lighting, locks, fire suppression systems, and burglar alarms.
- *Compensating controls* are controls designed to mitigate the risk associated with exceptions made to a security policy.

Exploring Compensating Controls

The Payment Card Industry Data Security Standard (PCI DSS) includes one of the most formal compensating control processes in use today. It sets out three criteria that must be met for a compensating control to be satisfactory:

- The control must meet the intent and rigor of the original requirement.
- The control must provide a similar level of defense as the original requirement, such that the compensating control sufficiently offsets the risk that the original PCI DSS requirement was designed to defend against.
- The control must be "above and beyond" other PCI DSS requirements.

For example, an organization might find that it needs to run an outdated version of an operating system on a specific machine because software necessary to run the business will only function on that operating system version. Most security policies would prohibit using the outdated operating system because it might be susceptible to security vulnerabilities. The organization could choose to run this system on an isolated network with either very little or no access to other systems as a compensating control.

The general idea is that a compensating control finds alternative means to achieve an objective when the organization cannot meet the original control requirement. Although PCI DSS offers a very formal process for compensating controls, the use of compensating controls is a common strategy in many different organizations, even those not subject to PCI DSS. Compensating controls balance the fact that it simply isn't possible to implement every required security control in every circumstance with the desire to manage risk to the greatest feasible degree.

In many cases, organizations adopt compensating controls to address a temporary exception to a security requirement. In those cases, the organization should also develop remediation plans designed to bring the organization back into compliance with the letter and intent of the original control.

Data Protection

Security professionals spend significant amounts of their time focusing on the protection of sensitive data. We serve as stewards and guardians, protecting the confidentiality, integrity, and availability of the sensitive data created by our organizations and entrusted to us by our customers and other stakeholders.

As we think through data protection techniques, it's helpful to consider the three states in which data might exist:

- *Data at rest* is stored data that resides on hard drives, tapes, in the cloud, or on other storage media. This data is prone to pilfering by insiders or external attackers who gain access to systems and are able to browse through their contents.
- *Data in motion* is data that is in transit over a network. When data travels on an untrusted network, it is open to eavesdropping attacks by anyone with access to those networks.
- *Data in processing* is data that is actively in use by a computer system. This includes the data stored in memory while processing takes place. An attacker with control of the system may be able to read the contents of memory and steal sensitive information.

We can use different security controls to safeguard data in all of these states, building a robust set of defenses that protects our organization's vital interests.

Data Encryption

Encryption technology uses mathematical algorithms to protect information from prying eyes, both while it is in transit over a network and while it resides on systems. Encrypted data is unintelligible to anyone who does not have access to the appropriate decryption key, making it safe to store and transmit encrypted data over otherwise insecure means.

We'll dive deeply into encryption tools and techniques in Chapter 7, "Cybersecurity Technology."

Data Loss Prevention

Data loss prevention (DLP) systems help organizations enforce information handling policies and procedures to prevent data loss and theft. They search systems for stores of sensitive information that might be unsecured and monitor network traffic for potential attempts to remove sensitive information from the organization. They can act quickly to block the transmission before damage is done and alert administrators to the attempted breach.

DLP systems work in two different environments:

- Host-based DLP
- Network DLP

Host-based DLP uses software agents installed on systems that search those systems for the presence of sensitive information. These searches often turn up Social Security numbers, credit card numbers, and other sensitive information in the most unlikely places!

Detecting the presence of stored sensitive information allows security professionals to take prompt action to either remove it or secure it with encryption. Taking the time to secure or remove information now may pay handsome rewards down the road if the device is lost, stolen, or compromised.

Host-based DLP can also monitor system configuration and user actions, blocking undesirable actions. For example, some organizations use host-based DLP to block users from accessing USB-based removable media devices that they might use to carry information out of the organization's secure environment.

Network-based DLP systems are dedicated devices that sit on the network and monitor outbound network traffic, watching for any transmissions that contain unencrypted sensitive information. They can then block those transmissions, preventing the unsecured loss of sensitive information.

DLP systems may simply block traffic that violates the organization's policy, or in some cases, they may automatically apply encryption to the content. This automatic encryption is commonly used with DLP systems that focus on email.

DLP systems also have two mechanisms of action:

- *Pattern matching*, where they watch for the telltale signs of sensitive information. For example, if they see a number that is formatted like a credit card or Social Security number, they can automatically trigger on that. Similarly, they may contain a database of sensitive terms, such as "credit card" or "blood pressure," and trigger when they see those terms in a transmission.
- *Watermarking*, where systems or administrators apply electronic tags to sensitive documents and then the DLP system can monitor systems and networks for unencrypted content containing those tags.

Watermarking technology is also commonly used in *digital rights management* (DRM) solutions that enforce copyright and data ownership restrictions.

Data Minimization

Data minimization techniques reduce risk by reducing the amount of sensitive information that we maintain on a regular basis. The best way to achieve data minimization is to simply destroy data when it is no longer necessary to meet our original business purpose.

If we can't completely remove data from a dataset, we can often transform it into a format where the original sensitive information is de-identified. The *de-identification* (or "anonymization") process removes the ability to link data back to an individual, reducing its sensitivity.

An alternative to de-identifying data is transforming it into a format where the original information can't be retrieved. This is a process called *data obfuscation*, and we have several tools at our disposal to assist with it:

- *Hashing* uses a hash function to transform a value in our dataset to a corresponding hash value. If we apply a strong hash function to a data element, we may replace the value in our file with the hashed value. Hashing uses a one-way function, meaning that it is not possible to retrieve the original value if you only have access to the hashed value.

- *Tokenization* replaces sensitive values with a unique identifier using a lookup table. For example, we might replace a widely known value, such as a student ID, with a randomly generated 10-digit number. We'd then maintain a lookup table that allows us to convert those back to student IDs if we need to determine someone's identity. Of course, if you use this approach, you must keep the lookup table secure!
- *Masking* partially redacts sensitive information by replacing some or all sensitive fields with blank characters. For example, we might replace all but the last four digits of a credit card number with X's or *'s to render the card number unreadable.

Although it isn't possible to retrieve the original value directly from the hashed value, there is one major flaw to this approach. If someone has a list of possible values for a field, they can conduct something called a *rainbow table attack*. In this attack, the attacker computes the hashes of those candidate values and then checks to see if those hashes exist in our data file.

For example, imagine that we have a file listing all the students at our college who have failed courses but we hash their student IDs. If an attacker has a list of all students, they can compute the hash values of all student IDs and then check to see which hash values are on the list. For this reason, hashing should only be used with caution.

Summary

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

As cybersecurity analysts seek to protect their organizations, they must evaluate risks to the CIA Triad. This includes the design and implementation of an appropriate mixture of security controls drawn from the managerial, operational, and technical control categories. These controls should also be varied in type, including a mixture of preventive, detective, corrective, deterrent, physical, and compensating controls.

Exam Essentials

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

Describe how information security strategies should be aligned with organizational goals and objectives. As information security managers develop their plans, they should use

reliable techniques to assess the current state of the program, such as threat research, SWOT analysis, and gap analysis. They may then identify the initiatives that will move the organization from the current state to its desired state.

Explain how security strategies are influenced by internal and external factors. Security strategies must be aligned with the business, but they must also incorporate other influences. Information security managers must remain abreast of emerging technologies, social media, the business environment, the organization's risk tolerance, regulatory requirements, third-party considerations, and the threat landscape as they develop, monitor, and revise cybersecurity strategies.

Know why stakeholder commitment and communication are essential to success. As information security leaders roll out new strategies, they must ensure that they have the support of senior leaders and other stakeholders. They may do this by clearly outlining how information security supports the organization's broader goals and objectives, identifying the business impact of security initiatives, and identifying clear success criteria.

Explain how security controls may be categorized based on their mechanism of action and their intent. Controls are grouped into the categories of managerial, operational, and technical based on the way that they achieve their objectives. They are divided into the types of preventive, detective, corrective, deterrent, compensating, and physical based on their intended purpose.

Describe the diverse impacts of data breaches on organizations. When an organization suffers a data breach, the resulting data loss often results in both direct and indirect damages. The organization suffers immediate financial repercussions due to the costs associated with the incident response, as well as long-term financial consequences due to reputational damage. This reputational damage may be difficult to quantify, but it may also have a lasting impact. In some cases, organizations may suffer operational damage if they experience availability damages, preventing them from accessing their own information.

Explain why data must be protected in transit, at rest, and in use. Attackers may attempt to eavesdrop on network transmissions containing sensitive information. This information is highly vulnerable when in transit unless protected by encryption technology. Attackers also might attempt to breach data stores, stealing data at rest. Encryption serves to protect stored data as well as data in transit. Data is also vulnerable while in use on a system and should be protected during data processing activities.

Know how data loss prevention (DLP) systems block data exfiltration attempts. DLP technology enforces information handling policies to prevent data loss and theft. DLP systems may function at the host level, using software agents to search systems for the presence of sensitive information. They may also work at the network level, watching for transmissions of unencrypted sensitive information. DLP systems detect sensitive information using pattern-matching technology and/or digital watermarking.

Explain how data minimization reduces risk by reducing the amount of sensitive information that we maintain. In cases where we cannot simply discard unnecessary information, we can protect information through de-identification and data obfuscation. The tools used to achieve these goals include hashing, tokenization, and masking of sensitive fields.

Review Questions

1. Matt is updating the organization's threat assessment process. What category of control is Matt implementing?
   - **A.** Operational
   - **B.** Technical
   - **C.** Corrective
   - **D.** Managerial

2. Jade's organization recently suffered a security breach that affected stored credit card data. Jade's primary concern is the fact that the organization is subject to sanctions for violating the provisions of the Payment Card Industry Data Security Standard. What category of risk is concerning Jade?
   - **A.** Strategic
   - **B.** Compliance
   - **C.** Operational
   - **D.** Financial

3. Chris is responding to a security incident that compromised one of his organization's web servers. He believes that the attackers defaced one or more pages on the website. What cybersecurity objective did this attack violate?
   - **A.** Confidentiality
   - **B.** Nonrepudiation
   - **C.** Integrity
   - **D.** Availability

4. Which one of the following elements is *most* important to gaining the support of senior leaders for cybersecurity initiatives?
   - **A.** Using plain, understandable language
   - **B.** Communicating often and in the format desired by the leaders
   - **C.** Demonstrating the alignment between business objectives and security needs
   - **D.** Adopting emerging technologies

5. Tonya is concerned about the risk that an attacker will attempt to gain access to her organization's database server. She is searching for a control that would discourage the attacker from attempting to gain access. What type of security control is she seeking to implement?
   - **A.** Preventive
   - **B.** Detective
   - **C.** Corrective
   - **D.** Deterrent

6. Which one of the following individuals bears ultimate responsibility for protecting an organization's data?
   - **A.** Data steward
   - **B.** End users
   - **C.** Data custodian
   - **D.** Data owner

7. Brooke is conducting a SWOT analysis for her organization's cybersecurity program. She recently learned about a cybersecurity insurance offering that may allow the organization to transfer some financial risk and is considering purchasing a policy. Where would this offering fit in the SWOT analysis?
   - **A.** Strength
   - **B.** Weakness
   - **C.** Opportunity
   - **D.** Threat

8. Tina is tuning her organization's intrusion prevention system to prevent false positive alerts. What type of control is Tina implementing?
   - **A.** Technical control
   - **B.** Physical control
   - **C.** Managerial control
   - **D.** Operational control

9. Dan is the CISO of an organization and he is spearheading the development of a new security operations center (SOC). He bears responsibility for the success of this initiative. In the RACI matrix entry for this initiative, how would Dan *best* be labeled?
   - **A.** R
   - **B.** A
   - **C.** C
   - **D.** I

10. Tony is reviewing the status of his organization's defenses against a breach of their file server. He believes that a compromise of the file server could reveal information that would prevent the company from continuing to do business. What term *best* describes the risk that Tony is considering?
    - **A.** Strategic
    - **B.** Reputational
    - **C.** Financial
    - **D.** Operational

11. Which one of the following data elements is not commonly associated with identity theft?

A. Social Security number

B. Driver's license number

C. Frequent flyer number

D. Passport number

12. What term best describes an organization's desired security state?

A. Control objectives

B. Security priorities

C. Strategic goals

D. Best practices

13. Jerry is developing a cybersecurity awareness program for members of his team who have administrative access to sensitive systems. What category *best* describes the users he is targeting?

A. Privileged users

B. High-risk users

C. End users

D. Data owners

14. Which one of the following individuals is the *least* appropriate direct manager of a chief information security officer?

A. Chief information officer

B. Chief risk officer

C. Chief executive officer

D. Senior director for identity and access management

15. Greg recently conducted an assessment of his organization's security controls and discovered a potential gap: the organization does not use full-disk encryption on laptops. What type of control gap exists in this case?

A. Detective

B. Corrective

C. Deterrent

D. Preventive

16. Toni is developing a new goal for her information security program. She has currently written it as "We will acquire and implement a new intrusion prevention system that will reduce successful network intrusions by 50%." What element of the SMART framework is lacking from this goal?

A. Specific

B. Measurable

C. Achievable

D. Relevant

E. Time-bound

17. Nolan is writing an after-action report on a security breach that took place in his organization. The attackers stole thousands of customer records from the organization's database. What cybersecurity principle was most impacted in this breach?

A. Availability

B. Nonrepudiation

C. Confidentiality

D. Integrity

18. Which one of the following objectives is not one of the three main objectives that information security professionals must achieve to protect their organizations against cybersecurity threats?

A. Integrity

B. Nonrepudiation

C. Availability

D. Confidentiality

19. What is the *most* appropriate span of control for a cybersecurity leader?

A. 2

B. 4

C. 7

D. 12

20. Brian is conducting a maturity assessment of his organization's cybersecurity team using Capability Maturity Model Integration (CMMI). He notes that the team does use defined processes but that they develop them in a reactive manner for each project they undertake. What level of maturity would *best* describe this team?

A. Defined

B. Repeatable

C. Initial

D. Quantitatively managed

E. Managed