DOMAIN 1

CISSP®

Security and Risk Management

DOMAIN 1 OF THE CISSP Common Body of Knowledge (CBK) covers the foundational topics of building and managing a risk-based information security program. This domain covers a wide variety of concepts upon which the remainder of the CBK builds.

Before diving into the heart of security and risk management concepts, this chapter begins with coverage of professional ethics and how they apply in the field of information security. Understanding your responsibilities as a security professional is equally as important as knowing how to apply the security concepts. We then move on to topics related to understanding your organization's mission, strategy, goals, and business objectives, and evaluating how to properly satisfy your organization's business needs securely.

Understanding risk management, and how its concepts apply to information security, is one of the most important things you should take away from this chapter. We describe risk management concepts and explain how to apply them within your organization's security program. In addition, understanding relevant legal, regulatory, and compliance requirements is a critical component of every information security program. Domain 1 includes coverage of concepts such as

cybercrimes and data breaches, import/export controls, and requirements for conducting various types of investigations.

This chapter introduces the human element of security and includes coverage of methods for educating your organization's employees on key security concepts. We cover the structure of a security awareness program and discuss how to evaluate the effectiveness of your education and training methods.

UNDERSTAND, ADHERE TO, AND PROMOTE PROFESSIONAL ETHICS

Understanding and following a strict code of ethics should be a top priority for any security professional. As a CISSP (or any information security professional who is certified by $(ISC)^2$), you are required to understand and fully commit to supporting the $(ISC)^2$ Code of Ethics. Any $(ISC)^2$ member who knowingly violates the $(ISC)^2$ Code of Ethics will be subject to peer review and potential penalties, which may include revocation of the member's $(ISC)^2$ certification(s).

$(ISC)^2$ Code of Professional Ethics

The $(ISC)^2$ Code of Ethics Preamble is as follows:

- The safety and welfare of society and the common good, duty to our principals, and to each other, requires that we adhere, and be seen to adhere, to the highest ethical standards of behavior.
- Therefore, strict adherence to this Code of Ethics is a condition of certification.

In short, the Code of Ethics Preamble states that it is required that every CISSP certified member not only follows the Code of Ethics but must be visibly seen as following the Code of Ethics. Even the perception of impropriety or ethical deviation may bring into question a member's standing. As such, CISSP certified members must serve as visible ethical leaders within their organizations and industry, at all times.

The $(ISC)^2$ Code of Ethics includes four canons that are intended to serve as high-level guidelines to augment, not replace, members' professional judgment. The $(ISC)^2$ Code of Ethics Canons are as follows:

- **Canon I:** Protect society, the common good, necessary public trust and confidence, and the infrastructure.
- **Canon II:** Act honorably, honestly, justly, responsibly, and legally.

- **Canon III:** Provide diligent and competent service to principals.
- **Canon IV:** Advance and protect the profession.

Adhering to and promoting the (ISC)² Code of Ethics not only includes being mindful of your own professional behaviors, but also being aware of your peers' behaviors. (ISC)² requires that any member who observes another member breaching the Code of Ethics follow the published ethics complaint procedure. Failure to do so may be considered breach of Canon IV. Additional information on the (ISC)² Code of Ethics and the ethics complaint procedures can be found at www.isc2.org/Ethics.

Organizational Code of Ethics

In addition to the (ISC)² Code of Ethics, as an information security professional, you must be aware of any code of ethics that you are required to uphold by your employer or industry. Similar to the (ISC)² Code of Ethics, these other organizational codes of ethics should not be considered replacements for sound judgment and moral behavior. As a CISSP, you are a leader within your organization. As such, you should lead by example in adhering to your organization's Code of Ethics.

✔ Ethics and the Internet

In January 1989, right around the dawn of the internet, the Internet Activities Board (IAB) released a memo titled "Ethics and the Internet" (RFC 1087) as a statement of policy concerning ethical use of the internet. Although the memo is ancient by technology standards, the principles within it are still relevant today; as a CISSP, you should understand and adhere to these principles.

RFC 1087 characterizes as unethical and unacceptable any activity that purposely

- "seeks to gain unauthorized access to the resources of the Internet"
- "disrupts the intended use of the Internet"
- "wastes resources (people, capacity, computer) through such actions
- "destroys the integrity of computer-based information
- "compromises the privacy of users"

Interestingly enough, this memo that debuted in the very early days of the internet — even before there was structured thought around information security — aligns directly with the CIA Triad and privacy principles that guide information security professionals in the 21st century.

UNDERSTAND AND APPLY SECURITY CONCEPTS

Information security refers to the processes and methodologies involved in safeguarding information and underlying systems from inappropriate access, use, modification, or disturbance. This is most often described by three critical security concepts: confidentiality, integrity, and availability. Together, these three principles form the pillars of information security known as the *CIA Triad* (see Figure 1.1).

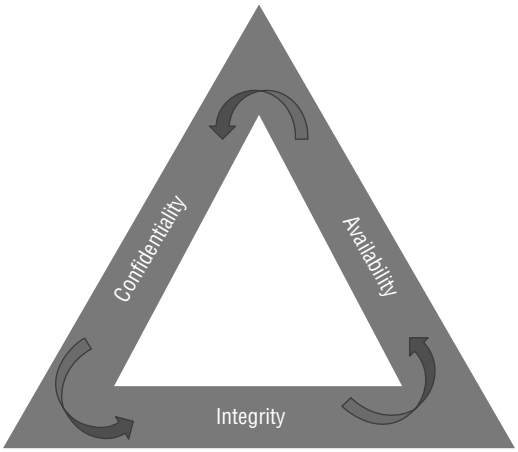

FIGURE 1.1 **CIA Triad**

Although different types of systems and data may prioritize one principle over the others, all three security concepts work together and depend on each other to successfully maintain information security. Confidentiality, integrity, and availability are the most critical characteristics that information security provides, and understanding each of these principles is a basic requirement for all information security professionals. As such, a common understanding of the meaning of each of the elements in the CIA Triad allows security professionals to communicate effectively.

Confidentiality

The first principle of the CIA Triad is confidentiality. *Confidentiality* is the concept of limiting access to data to authorized users and systems and restricting access from unauthorized parties. In other words, confidentiality guarantees that only intended people are able to access information and resources. In short, the goal of confidentiality is to prevent unauthorized access to data, as it pertains to viewing, copying, or moving the data. Organizations that fail to protect data confidentiality run the risk of violating contractual

and regulatory obligations around data protection. In addition, companies with known breaches of confidentiality commonly experience reputational damage that benefits their competitors and potentially impacts their bottom line.

The concept of confidentiality is closely related to the security best practice of *least privilege*, which asserts that access to information should be granted only on a need-to-know basis. Under least privilege, users are only given enough access to do their jobs — no more and no less.

Privacy is a field that is closely related to security and is focused on the confidentiality of personal data specifically. Information that can be used to uniquely identify a person, such as names, birthdates, and Social Security numbers, are considered *personally identifiable information* (PII), which is usually required by law to be kept highly confidential. Privacy requirements are introduced later in this chapter, in the section "Determine Compliance and Other Requirements."

Numerous malicious acts target data confidentiality, including phishing and other forms of social engineering, credential (e.g., password) theft, network sniffing, and others. In addition to malicious acts, confidentiality may be compromised by human error, oversight, or mere negligence. Such examples include failure to encrypt sensitive data, misrouted emails, or displaying sensitive information on your computer monitor while unauthorized viewers are in the vicinity.

Confidentiality is often the security concept that data owners care about the most, and there are many security controls available to assist with this. Encryption, multifactor authentication, and role-based access controls are a few measures that can help ensure data confidentiality. Extensive personnel training is a hugely important measure for reducing risk associated with human error and negligence. We address data confidentiality and discuss relevant controls throughout the remainder of this book.

Integrity

The second principle of the CIA Triad is integrity. *Integrity* is the concept of maintaining the accuracy, validity, and completeness of data and systems. It ensures that data is not manipulated by anyone other than an authorized party with an authorized purpose, and that any unauthorized manipulation is easily identifiable as such. The primary goal of integrity is to ensure that all data remains intact, correct, and reliable. Failure to properly protect data integrity can have a negative impact on business processes, including leading to personnel making improper decisions or potentially harmful actions, due to having incorrect information.

As with confidentiality, integrity may be compromised by malicious acts, human error, oversight, or negligence. Viruses, compromise of poorly written code, and intentional modification are examples of malicious threats that may violate integrity, and we discuss others throughout this book. In addition, integrity violations can occur when an

administrator enters the wrong command in a database, when an administrator alters the wrong line in a configuration file, or when a user accidentally introduces malware into their system through a phishing email.

Data backups, software version control, strict access control, and cryptographic hashes are some measures that help ensure data integrity by preventing unauthorized modifications or by allowing tampered data to be restored to a known-good state. Similar to confidentiality, extensive security awareness training is a major factor in preventing nonmalicious integrity violations.

Authenticity and nonrepudiation are two concepts that are closely related to integrity. *Authenticity* refers to ensuring that data is genuine and that all parties are who they say they are. *Nonrepudiation* is a legal principle that has a strong parallel in the information security world; this concept requires ensuring that no party is able to deny their actions (e.g., creating, modifying, or deleting data). Digital signatures are the most common mechanisms used to establish authenticity and nonrepudiation in information systems.

Availability

The third and final principle of the CIA Triad is availability. *Availability* is the concept focused on ensuring that authorized users can access data when they need it. In enterprise environments, the concept of availability entails providing assurance that legitimate parties have timely and uninterrupted access to the systems and data that they need to do their jobs. Threats against availability can interfere or even halt an organization's business operations. An extended disruption of critical systems and data may lead to reputational damage that results in loss of customers and revenue.

Related concepts that should be considered alongside availability include the following:

- *Accessibility* refers to the ability and ease of a user to use a resource or access data when needed. This involves removing barriers for authorized users to access these resources and data. For example, consider a file that's stored on your company's internal network drive. As long as the file is intact and the network drive is up and running, that file can be considered available. However, if someone were to move that file to a protected folder on the shared drive, you may lack the required permissions to access that file — the data is still available but is no longer accessible.
- *Usability* refers to the ability of a user to meet their needs with available data. If you have ever needed to edit a Google doc (or any other file) and noticed that you have been granted only read-only permissions, then that file was absolutely available but lacked sufficient usability.
- *Timeliness* refers to the time expectation for availability of information and resources and is the measure of the time between when information is expected

and when it is available for use. Ensuring timeliness requires that data is available to authorized users within an acceptable period of time. For cloud services and other situations that involve a third party managing data, timeliness is a key factor that must be agreed upon and documented in service level agreements (SLAs).

There are many threats to data and system availability, and they may be either malicious or nonmalicious, either man-made or naturally occurring. Malicious availability threats include denial-of-service (DoS) attacks, object deletion, and ransomware attacks. While malicious compromise of availability tends to get all the buzz, there are various nonmalicious threats that can interfere with resource and data availability. Some common examples include hardware failures, software errors, and environmental threats such as power outages, floods, excessive heat, and so forth. When planning your information security program, it's essential that you thoroughly consider both human-based and naturally occurring threats and develop mitigations that address all threat vectors.

Mechanisms such as data backups, redundant storage, backup power supply, and web application firewalls (WAFs) can help prevent disruption of system and information availability. For systems that have a requirement for high availability and continuous uptime, cloud computing offers added redundancy and extra assurance of availability.

Limitations of the CIA Triad

The CIA Triad evolved out of theoretical work done in the mid-1960s. Precisely because of its simplicity, the rise of distributed systems and a vast number of new applications for new technology has caused researchers and security practitioners to extend the triad's coverage.

Guaranteeing the identities of parties involved in communications is essential to confidentiality. The CIA Triad does not directly address the issues of authenticity and nonrepudiation, but the point of nonrepudiation is that neither party can deny that they participated in the communication. This extension of the triad uniquely addresses aspects of confidentiality and integrity that were never considered in the early theoretical work.

The National Institute of Standards and Technology (NIST) Special Publication 800-33, "Underlying Technical Models for Information Technology Security," included the CIA Triad as three of its five security objectives, but added the concepts of accountability (that actions of an entity may be traced uniquely to that entity) and assurance (the basis for confidence that the security measures, both technical and operational, work as intended to protect the system and the information it processes). The NIST work remains influential as an effort to codify best-practice approaches to systems security.

Perhaps the most widely accepted extension to the CIA Triad was proposed by information security pioneer Donn B. Parker. In extending the triad, Parker incorporated

three additional concepts into the model, arguing that these concepts were both atomic (could not be further broken down conceptually) and nonoverlapping. This framework has come to be known as the Parkerian Hexad. The Parkerian Hexad contains the following concepts:

- **Confidentiality:** The limits on who has access to information
- **Integrity:** Whether the information is in its intended state
- **Availability:** Whether the information can be accessed in a timely manner
- **Authenticity:** The proper attribution of the person who created the information
- **Utility:** The usefulness of the information
- **Possession or control:** The physical state where the information is maintained

Subsequent academic work produced dozens of other information security models, all aimed at the same fundamental issue — how to characterize information security risks.

In addition to security topics codified in the CIA Triad and related models, the concept of privacy has grown to be a core consideration of security professionals. *Privacy*, as defined in the (ISC)^2 glossary, is the right of human individuals to control the distribution of information about themselves. Privacy, though often managed outside of organizations' central security team, is closely related to the principle of confidentiality and must be a priority for every organization that handles employee or customer personal information. We discuss privacy in several sections throughout the rest of this book.

For the security professional, a solid understanding of the CIA Triad is essential when communicating about information security practice, but it's important to consider related topics not covered by the triad.

EVALUATE AND APPLY SECURITY GOVERNANCE PRINCIPLES

Security governance is the set of responsibilities, policies, and procedures related to defining, managing, and overseeing security practices at an organization. Security is often mistakenly considered to be an IT issue; in actuality, securing an organization's assets and data is a business issue and requires a high level of planning and oversight by people throughout the entire organization, not just the IT department. Because security is a wide-ranging business issue, security governance commonly overlaps with corporate governance and IT governance for an organization. As such, security governance is typically led by executive management at a company, usually including the board of directors. Applying security governance principles involves the following:

- Aligning the organization's security function to the company's business strategy, goals, mission, and objectives

- Defining and managing organizational processes that require security involvement or oversight (e.g., acquisitions, divestitures, and governance committees)
- Developing security roles and responsibilities throughout the organization
- Identifying one or more security control frameworks to align your organization with
- Conducting due diligence and due care activities on an ongoing basis

Alignment of the Security Function to Business Strategy, Goals, Mission, and Objectives

An effective security function must be in alignment with the company's business strategy, goals, mission, and business objectives. Each of these elements should be considered during the creation and management of the organization's information security program and policies.

Companies that fail to properly align their security program with their business strategy, goals, mission, and objectives often perceive security as a business blocker; these companies frequently experience information security as a hurdle that must be cleared to get things accomplished. On the contrary, an information security function that is tightly aligned with a company's strategy and mission can serve as a business enabler, where security is built into the fabric of the company and helps drive toward common goals and objectives. In other words, a company should achieve its mission thanks in part to security, not *despite* security.

A *mission statement* is a simple declaration that defines a company's function and purpose; a mission statement summarizes what the company is, what it does, and why the company exists to do those things. A mission statement should be used to drive all corporate activities, including the organization's allocation of time, finances, and effort.

A *business strategy* describes the actions that a company takes to achieve its goals and objectives. Whereas a mission statement describes *what* will be achieved, an organization's business strategy identifies exactly *how* the mission will be accomplished. A company's mission statement rarely changes, but an organization's strategy must be flexible enough to change as the business environment changes.

A *goal*, in business, is something that an organization expects to achieve or accomplish. Business goals help a company plan for success, and an organization's goals should contribute to its mission. Many companies use the SMART criteria to define their organizational goals. *SMART* is a mnemonic acronym that defines criteria for creating quality goals. A SMART goal must exhibit the following characteristics:

- **Specific:** State what you will do using real numbers and real deadlines.
- **Measurable:** Identify a way to evaluate progress and measure success (or failure). Use metrics or data targets to ensure that the goal is trackable.

- **Achievable or Attainable:** Establish challenging, but possible, goals that are within your scope.
- **Relevant:** Establish a goal that is pertinent to your overall mission and vision and aligned with your organization's values and strategy.
- **Time-bound:** State when you will get the goal done, using specific dates or timeframes.

An *objective* is a milestone or a specific step that contributes to an organization reaching its goals and achieving its mission. Objectives are used to define incremental steps toward achieving a broader goal. Much like SMART goals, organizations often use the SMART framework to define quality objectives. While many people incorrectly use the terms *goal* and *objective* interchangeably, you should understand that an objective is a short-term milestone that supports a longer-term goal.

When establishing your organization's security function, you should begin by defining a security strategy that aligns with your organization's overall business strategy and mission statement. You should develop a set of specific, measurable, achievable, relevant, and time-bound goals and objectives that will help you efficiently maintain the confidentiality, integrity, and availability of your company's systems and information without disrupting your organization's ability to achieve its business goals and objectives. Running an effective security program demands careful consideration of business needs and organizational strategy, in addition to legal and compliance requirements, and requires governance to manage the effectiveness of the security function within the overall organization.

Organizational Processes

People who consider information security a purely IT matter are more prone to focusing solely on the technologies that fit into a security program. As a CISSP, you should know that a mature information security program is more than a collection of firewalls, intrusion detection systems and intrusion prevention systems (IDSs/IPS), and other tools thrown together — a well-managed security program requires processes in place to provide oversight of activities by executive members of the organization. *Security governance* is the set of all organizational processes involved in defining and managing information security policies and procedures, including the oversight to ensure that those policies and procedures follow the direction of the organization's strategy and mission.

Governance Committees

A *governance committee* is a group of executives and leaders who regularly meet to set the direction of the company's security function and provide guidance to help the security function align with the company's overall mission and business strategy. Governance committees review ongoing and planned projects, operational metrics, and any other security matters that may concern the business as a whole. The primary objective of a governance committee is to provide oversight for the company's security function, while ensuring that the security function continues to meet the needs of the organization and its stakeholders.

There are many organizational processes that require a heavy dose of security governance. Mergers, acquisitions, and divestitures are major business events that come with a great deal of security risk that a company must manage.

Mergers and Acquisitions

A *merger* is the combining of two separate organizations that creates a new, joint organization. An *acquisition* is the takeover of one organization by another. While mergers and acquisitions (M&A) have different business approaches, they share many of the same security concerns and are often discussed together.

There are countless potential security risks when a company acquires another company or when two organizations decide to merge. For any merger or acquisition, it's imperative that organizations consider these risks and identify appropriate mitigations before pulling the trigger. Some M&A risk factors to consider include the following:

- **Absorbing the unknown:** When merging with or acquiring another organization, you are absorbing its entire IT infrastructure — good or bad. This means that you are acquiring systems that are likely managed differently from your own, and there may be significant differences in the security controls and processes in place. In addition, the acquired company may use homegrown or highly customized applications that will need to be securely integrated into your existing environment. Further, the acquired or absorbed company may use a different approach to threat modeling and vulnerability management (if they do these at all). Differences in security processes may result in operational challenges and inconsistent procedures during and after integration of the two businesses.
- **Creating new attack vectors:** By adding in new systems and platforms, you are potentially creating new routes for your company to be attacked. For example, if your organization uses Windows and macOS and you acquire a company that has a fleet of Linux systems, you now have a third operating system to manage and secure.

- **Impacting resources:** Mergers and acquisitions are challenging for everyone involved. Your IT teams may be stretched thin as they try to come up to speed on the newly acquired infrastructure while also keeping the existing systems running securely.
- **Disgruntled employees:** In addition to the burnout that employees may feel, corporate M&A can cause severe dissatisfaction in employees who were completely happy in their previously standalone company. Insider threat is commonly considered a top security concern, and acquiring disgruntled employees poses a severe threat to an organization.

A company's security function should play a key role during any M&A conversations and should help identify potential issues related to the target organization's security posture, operations, or policy. If the two organizations have greatly different security postures, it's usually best to reconsider the deal or, at the least, consider fixing key security gaps before connecting the two company's networks and systems. Considering the risks from the previous list, an acquiring company (or both companies in a merger) should perform the following activities prior to completing an M&A deal:

- Review the company's information security policies and procedures to determine how thorough they are. Take note of any missing security policies/procedures that may indicate a low level of security maturity.
- Review the company's data assets and determine any applicable regulatory or compliance requirements (e.g., PCI, HIPAA, GDPR, etc.). Take particular note of any regulations that may present new requirements for your organization.
- Review the organization's personnel security policies to identify potential issues or concerns. For example, your company may have compliance requirements to conduct a specific type of background check, and the target company may not be compliant.
- Identify any proprietary or custom applications managed by the company and request static and dynamic application security tests be run against them to demonstrate their security posture. (SAST and DAST are covered in Chapter 8, "Software Development and Security.")
- Request results from a recent penetration test (pentest) that includes network, operating system, application, and database testing. Any critical or high findings should have a plan for remediation or mitigation.
- Review the organization's use of third-party and open-source software to ensure that the software is safe and appropriately licensed.

The previous list is not intended to be comprehensive, but rather a starting point for things to consider prior to any mergers and acquisitions. Your security team needs to be an integral part of the M&A process from the initial talks through integration.

Divestitures

A *divestiture* is the act of selling off or disposing of a subset of business interests or assets. An organization may pursue a divestiture for a number of reasons: political, social, or strictly financial. Divestitures often occur when management decides that a certain part of the business no longer aligns with the company's business strategy or mission. Divestitures also frequently happen after a merger or acquisition, in cases where the merger or acquisition creates redundancies within the combined organization.

Information usually accompanies the physical assets and interests that a company divests, which presents a major concern for information security. The biggest security concern in a divestiture involves maintaining confidentiality as the company gets rid of assets that may contain sensitive or proprietary information. As a CISSP, you should ensure that your organization takes the following actions prior to completing a divestiture:

- Identify and categorize all assets that are involved in the divestiture; this includes hardware, software, and information assets. Creating a complete inventory of all impacted assets is a critical first step to ensuring a secure divestiture.
- Decouple impacted systems from your remaining infrastructure. Your company likely uses common human resources (HR), accounting, and technology systems (such as a virtual private network, email, etc.) to support the entire company. The assets being divested must be removed from this common infrastructure and spun out for the new organization to own and manage.
- Review all access permissions. You must identify who has access to the impacted assets and determine whether they need to maintain that access. People are sometimes part of a divestiture, and a subset of the employee base may leave with other divested assets. If that is the case in your divestiture, you must appropriately revoke unnecessary permissions while leaving required permissions intact.
- Consult your legal and compliance teams to ensure that you follow all required regulatory and compliance requirements around data retention, deletion, etc.

During a divestiture, both companies (i.e., the divesting and the divested company) must consider how the business transaction impacts their respective security program. Each company must ensure that their security controls, operations, policies and procedures, and governance structure continue to support the newly restructured companies. If the divested company was sold to another company (i.e., as part of an acquisition), then the purchasing company must update its security program to accommodate for its newly acquired assets. In cases where a divested company leads to the formation of a completely new entity, the new company must create an all-new security function (and the supporting policies, procedures, and governance structure) to appropriately manage information security.

Much like mergers and acquisitions, divestitures can present a number of security challenges for organizations. Similarly, the key to a successful divestiture is active involvement by your security team from the early planning phases all the way through completion of the divestiture.

Organizational Roles and Responsibilities

People who don't work in security often look at security professionals as the only employees responsible for keeping systems and information secure. Of course, as information security professionals, we know that it is really *everyone's* job to keep the organization's assets and information secure — that means from the chief executive officer (CEO) down to the most junior clerk in the mailroom, and everyone in between. As a CISSP, one of your jobs is to evangelize security throughout your company, while helping to define security roles and responsibilities throughout the organization.

Your organization should define security roles and responsibilities within your information security policy, and it should align with roles and responsibilities defined in other organizational policies. It's important that roles and responsibilities are defined for and understood by employees of every level and line of business, as well as third parties such as contractors and vendors.

Different companies have different roles, but the following designations are some of the most commonly seen information security roles:

- **Chief information security officer (CISO):** A *CISO* is the senior-level executive within an organization who is responsible for the overall management and supervision of the information security program. The CISO drives the organization's security strategy and vision and is ultimately responsible for the security of the company's systems and information. While corporate reporting structures vary by company size and industry, most CISOs now report to a company's chief information officer (CIO) or CEO.
- **Chief security officer (CSO):** A *CSO* is a senior-level executive within an organization who is generally responsible for all physical security and personnel security matters. Many organizations have merged CSO responsibilities into the CISO role, but you should be aware of the potential distinction between the two. To make matters even more confusing, some organizations refer to their overall security leader as a CSO (instead of CISO). You should lean on context anytime you see these titles used.
- **Security analyst:** A *security analyst* is someone with technical expertise in one or more security domains who executes the day-to-day security work. This may include things such as data analysis, firewall rule management, incident handling, and other operational activities.

- **Manager or program manager:** In security, a *manager* (or *program manager*) is someone who owns one or more processes related to information security. A security manager may be the owner for compliance, vulnerability management, or any other broad set of responsibilities that are executed by security analysts.
- **Director:** In security, a *director* is generally a manager of managers who is responsible for the overall strategic guidance of a group of security programs.

NOTE While the role of CISO has traditionally reported to a company's CIO, that trend is changing. Organizations increasingly view information security as not only an IT issue but a business issue. As a result, many argue that CISOs should report directly to a company's CEO.

As previously mentioned, security is everyone's responsibility. Outside of the information security roles and responsibilities described in the previous list, every user within an organization plays an important role in keeping information secure. A *user* (or *end user*) includes any person who accesses or handles an organization's information systems or data. Users may include full-time and part-time employees, interns, contractors, consultants, vendors, partners, and so on. Some general user responsibilities include the following:

- Understand, agree to, and adhere to all information security policies, procedures, standards, and guidelines, as well as any relevant regulatory and compliance requirements. Users are also responsible for satisfying contractual obligations (such as nondisclosure agreements) that affect the confidentiality of the company's information and processes.
- Complete all required information security training and awareness activities by their required completion dates.
- Report any actual or suspected security violations or breaches to appropriate personnel in a timely manner.

Security Control Frameworks

Poor security management is one of the primary culprits for many organizations' security problems. Security management can be accomplished by adopting a top-down approach, bottom-up approach, or some combination of the two.

Historically, enterprises have utilized more of a *bottom-up approach* to security, in which the IT department takes security seriously and attempts to develop a security function for the company. With this approach, operations staff identify security needs and issues and push those findings up the chain to senior management to provide guidance and funding. This approach tends to result in a company taking reactive

measures rather than instituting proactive policies, and often leads to underfunded security programs.

In a *top-down approach*, senior leadership starts by understanding the regulations and security threats faced by the organization, and initiates strategies, policies, and guidelines that are pushed down throughout the rest of the organization. With a top-down approach, information security is evangelized by the most senior executives at the company, which ensures that security is prioritized and in alignment with the company's overall business strategy. An effective top-down approach requires strong governance (as discussed earlier in this chapter) that starts with aligning with one or more security control frameworks.

A *security control* is a technical, operational, or management safeguard used to prevent, detect, minimize, or counteract security threats. (ISC)^2 defines a *security control framework* as "a notional construct outlining the organization's approach to security, including a list of specific security processes, procedures, and solutions used by the organization." Organizations often adopt a security control framework to assist with meeting legal and regulatory compliance obligations, while also helping to build a security program that maintains the confidentiality, integrity, and availability of the company's assets.

NOTE Technical controls are system-based safeguards and countermeasures — things like firewalls, IDS/IPS, and data loss prevention (DLP). Operational controls are safeguards and countermeasures that are primarily implemented and executed by people (as opposed to systems); security guards are a common example. Management controls include policies, procedures, and other countermeasures that control (or manage) the information security risk. Management controls are sometimes referred to as *administrative controls*, but this should not be confused with activities associated with a system admin (sysadmin). The lines between the three categories can often blur (i.e., many controls fit into more than one of the categories), and many organizations have discontinued use of the terms to avoid confusion. You should be familiar with the concepts, should you come across the terminology in your organization.

Organizations often select security control frameworks based on their industry. For example, the Payment Card Industry (PCI) control framework is a global framework used by financial services organizations and companies that handle cardholder data. In the United States, the Health Insurance Portability and Accountability Act (HIPAA) offers a control framework for healthcare organizations and companies that handle protected health information (PHI). Aside from PCI and HIPAA (which are covered later in this chapter), ISO/IEC, NIST, and CIS provide some of the most frequently adopted security control frameworks used by organizations across multiple industries. While there

are many other control frameworks available, you should be familiar with the following frameworks, at a minimum:

- ISO/IEC 27001
- ISO/IEC 27002
- NIST 800-53
- NIST Cybersecurity Framework
- CIS Critical Security Controls

NOTE The HITRUST (originally known as the Health Information Trust Alliance) Common Security Framework (CSF) was originally developed to address the overlapping regulatory environment in which many healthcare providers operate. It has evolved over time to provide a comprehensive, prescriptive framework that can be used for organizations that exchange any type of sensitive and/or regulated data. Taking into account both risk-based and compliance-based considerations, the HITRUST provides an auditable framework for the evaluation of an organization's security environment.

NOTE Control Objectives for Information Technologies (COBIT) is a framework developed by ISACA (previously known as the Information Systems Audit and Control Association) for overall information technology management and governance and is perhaps the most popular IT governance framework used in industry. While it is not a security-specific control framework, it does outline end-to-end IT governance objectives and processes that encompass many security requirements and concepts. Visit www.isaca.org if you want to learn more about the COBIT framework.

ISO/IEC 27001

ISO/IEC 27001 (sometimes referred to as just ISO 27001) is an information security standard published by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). ISO 27001 is the most popular standard within the ISO/IEC 27000 family of standards and is focused on the creation and maintenance of an *information security management system* (ISMS), which ISO defines as "a systematic approach to managing sensitive company information so that it remains secure." In plain English, an ISMS is a set of people, processes, and technologies that manages the overall security of a company's systems and data. ISO/IEC 27001 describes

the overall components of an ISMS, and this standard is the basis for many organization's security programs.

As of this writing, the most recent revision to ISO/IEC 27001 was in 2013, though its parent, ISO/IEC 27000, was revised in 2018. ISO 27001:2013 contains 114 controls across 14 domains, as follows:

- Information security policies
- Organization of information security
- Human resource security
- Asset management
- Access control
- Cryptography
- Physical and environmental security
- Operations security
- Communications security
- System acquisition, development, and maintenance
- Supplier relationships
- Information security incident management
- Information security aspects of business continuity management
- Compliance

ISO/IEC 27002

ISO/IEC 27002 (again, often referred to as just ISO 27002) is titled "Security Techniques — Code of practice for information security controls." This standard builds on ISO 27001 by providing guidelines for organizations to select, implement, and manage security controls based on their own security risk profile. In other words, ISO 27002 is a bit more prescriptive than ISO 27001, as it provides best-practice recommendations for organizations to build and maintain their ISMSs.

NIST 800-53

The National Institute of Standards and Technology is a nonregulatory agency of the U.S. Department of Commerce, whose mission is to promote innovation and industrial competitiveness by advancing standards and technologies. NIST publishes and manages a variety of special publications related to information security, cloud computing, and other technologies. NIST 800-53, "Security and Privacy Controls for Federal Information Systems and Organizations," is NIST's massive security control framework. Though NIST

800-53 was initially created to aid U.S. government agencies in managing their security programs, it is widely regarded as one of the most comprehensive baselines of security controls and is referenced across many industries around the globe. NIST 800-53 defines hundreds of security controls across the following 18 control families:

- Access control (AC)
- Awareness and training (AT)
- Audit and accountability (AU)
- Security assessment and authorization (CA)
- Configuration management (CM)
- Contingency planning (CP)
- Identification and authentication (IA)
- Incident response (IR)
- Maintenance (MA)
- Media protection (MP)
- Physical and environmental protection (PE)
- Planning (PL)
- Personnel security (PS)
- Risk assessment (RA)
- System and services acquisition (SA)
- System and communications protection (SC)
- System and information integrity (SI)
- Program management (PM)

NOTE The latest revision of NIST 800-53, Rev. 5, was released in September 2020.

NIST Cybersecurity Framework

The NIST Cybersecurity Framework (CSF), first published in 2014, is a collection of standards, guidelines, and best practices to manage cybersecurity risk. As of this writing, NIST CSF v1.1 is the current version and was released in 2018. NIST CSF was initially developed with a focus on industries considered "critical infrastructure" — industries such as banking, energy, and communications. It has since become a go-to controls framework for companies of all sizes and across all business sectors.

The NIST CSF aligns with controls and best practices in NIST 800-53 and other control frameworks, but was designed to be a more flexible and understandable option for private-sector companies to adapt. The NIST Cybersecurity Framework consists of five core functions, each with multiple subdivisions NIST calls categories. (See Figure 1.2.)

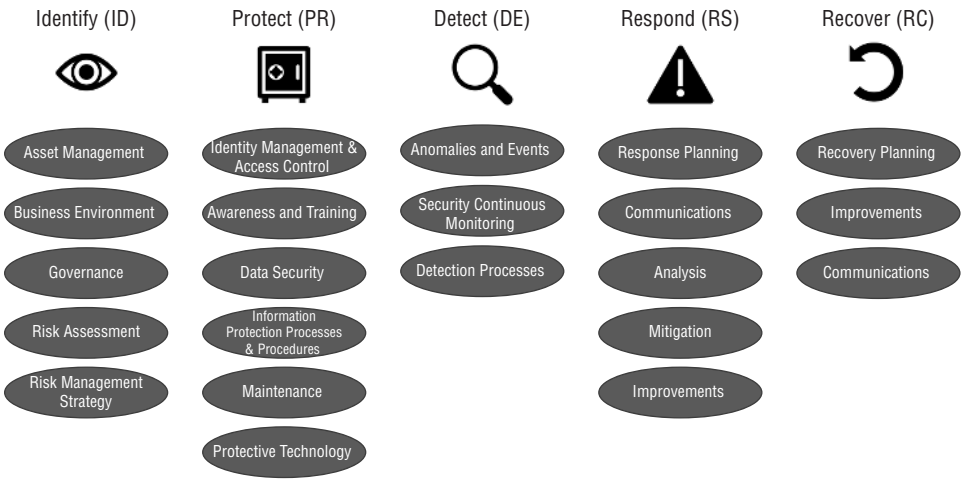

FIGURE 1.2 **NIST Cybersecurity Framework**

The five core functions within NIST CSF are intended to be performed simultaneously and continuously to form a culture of assessing and addressing cybersecurity risk. NIST defines the purpose of each of the five core functions as follows:

- **Identify (ID):** Develop an organizational understanding to manage cybersecurity risk to systems, people, assets, data, and capabilities.
- **Protect (PR):** Develop and implement appropriate safeguards to ensure delivery of critical services.
- **Detect (DE):** Develop and implement appropriate activities to identify the occurrence of a cybersecurity event.
- **Respond (RS):** Develop and implement appropriate activities to take action regarding a detected cybersecurity incident.
- **Recover (RC):** Develop and implement appropriate activities to maintain plans for resilience and restore any capabilities or services that were impaired due to a cybersecurity incident.

The five core functions are divided into 23 categories, and these categories are further divided into a total of 108 subcategories. Each subcategory describes a specific security

control or desired outcome. Visit www.nist.gov/cyberframework for the complete list of subcategories and additional guidance on using the NIST Cybersecurity Framework.

CIS Critical Security Controls

The CIS Critical Security Controls (or CIS Controls) is a publication of 20 best-practice guidelines for information security. The publication was initially created by SANS Institute but was transferred to the Center for Internet Security (CIS) in 2015. Today, you may see these 20 critical controls labeled CIS CSC, CIS 20, Sans Top 20, or other variants.

CIS Controls v7.1 was released in April 2019, and identifies the basic, foundational, and organizational controls that CIS recommends mitigating the most common attacks against networks and systems. According to the Center for Internet Security, the 20 Critical Security Controls are as follows:

- **CIS Control 1:** Inventory and Control of Hardware Assets
- **CIS Control 2:** Inventory and Control of Software Assets
- **CIS Control 3:** Continuous Vulnerability Management
- **CIS Control 4:** Controlled Use of Administrative Privileges
- **CIS Control 5:** Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
- **CIS Control 6:** Maintenance, Monitoring, and Analysis of Audit Logs
- **CIS Control 7:** Email and Web Browser Protections
- **CIS Control 8:** Malware Defenses
- **CIS Control 9:** Limitation and Control of Network Ports, Protocols, and Services
- **CIS Control 10:** Data Recovery Capabilities
- **CIS Control 11:** Secure Configuration for Network Devices, such as Firewalls, Routers, and Switches
- **CIS Control 12:** Boundary Defense
- **CIS Control 13:** Data Protection
- **CIS Control 14:** Controlled Access Based on the Need to Know
- **CIS Control 15:** Wireless Access Control
- **CIS Control 16:** Account Monitoring and Control
- **CIS Control 17:** Implement a Security Awareness and Training Program
- **CIS Control 18:** Application Software Security
- **CIS Control 19:** Incident Response and Management
- **CIS Control 20:** Penetration Tests and Red Team Exercises

NOTE The controls and subcontrols within the CIS CSC break down into what are known as Implementation Groups. According to CIS, "Implementation Groups provide a simple and accessible way to help organizations of different classes focus their security resources, and still leverage the value of the CIS Controls program . . . ." In essence, these Implementation Groups help organizations prioritize controls and identify the subcontrols that are most reasonable for level of expertise and their risk profile. Visit www.cissecurity.org for more information on the CSC and their Implementation Groups.

Due Care and Due Diligence

Governance requires that the individuals setting the strategic direction and mission of the organization act on behalf of the stakeholders. The minimum standard for their governance action requires that they act with due care. *Due care* is a legal term used to describe the conduct that a reasonable person would exercise in a given situation. In business, due care is using reasonable care to protect the interests of your organization. More specifically, in regard to information security, due care relates to the conduct that a reasonable person would exercise to maintain the confidentiality, integrity, and availability of their organization's assets. This concept of "reasonable" can be a bit nebulous at first, but it is intended to protect a person or organization from accusations of negligence. In short, court decisions around the world have demonstrated that a person's actions can be assumed "reasonable" if a person of similar background and experience, confronted with the same situation would enact the same or similar actions. Examples of due care in security are activities like scanning and patching security vulnerabilities, enabling security logging, and writing restrictive firewall rules that enforce least privilege (discussed in Chapter 3, "Security Architecture and Engineering").

Due diligence is another legal concept that relates to continually ensuring that behavior maintains due care. In other words, due diligence is the ongoing execution and monitoring of due care. In relation to information security, due diligence relates to the ongoing actions that an organization and its personnel conduct to ensure organizational assets are reasonably protected. Examples of due diligence in security are activities like reviewing security log output for suspicious activity and conducting penetration tests to determine if firewall rules are sufficiently restrictive.

The concepts of due care and due diligence are incredibly important in the legal and finance world, but they must also be understood by information security professionals. Exercising due care and conducting due diligence are required to avoid claims of negligence in court. The CISSP CBK aims to establish the set of knowledge and activities required of a "reasonable" security leader.

DETERMINE COMPLIANCE AND OTHER REQUIREMENTS

(ISC)^2 defines *compliance* as adherence to a mandate; it includes the set of activities that an organization conducts to understand and satisfy all applicable laws, regulatory requirements, industry standards, and contractual agreements.

Legislative and Regulatory Requirements

Many compliance expectations come from statutory or regulatory requirements that apply broadly to all industries. Others are specific to certain industries or products. This ever-changing set of expectations requires a continuous review of organizational practices to ensure that information is protected in compliance with all applicable requirements.

NOTE Because there are many compliance requirements that relate to information security, many people often confuse the two or assume that being compliant is the same as being secure. As a CISSP, you should understand that compliance requirements generally serve as a solid *baseline* for security, but being compliant with security regulations and standards is only the first step toward being secure.

The first challenge in identifying compliance requirements involves knowing which jurisdiction has the legal authority to set those requirements. *Jurisdiction* is a legal concept that establishes the official power to make legal decisions and judgments. It is not enough to know the relevant geography or political boundaries; jurisdiction may also be influenced by international treaties and agreements, the activity of your organization, or any number of other factors. Regardless of the example laws and regulations listed in this text, information security practitioners must be aware of the nuances of the jurisdictions in which they operate.

In most jurisdictions, laws are established to define what is permissible and what is not. In U.S. law, the word *law* refers to any rule that, if broken, subjects a party to criminal punishment or civil liability. Laws may be generally categorized into two parts: statutes and regulations. Statutes are written and adopted by the jurisdiction's legislative body (e.g., U.S. Congress), while regulations are more detailed rules on how the execution of a statute will be performed. Both statutes and regulations are legally enforceable, but regulations are subordinate to statutes.

TIP While you don't need a law degree to be an information security professional, it certainly helps to Article 17 have a basic understanding of legal jargon. Since you must learn to read laws and understand how they apply to information security, it's important that you can interpret how laws are usually cited, particularly in the United States. If you see 18 U.S.C. § 1030, for example, you should understand that this refers to Section 1030 of Title 18 of the United States Code. You may see C.F.R. used to reference the Code of Federal Regulations. In the United Kingdom, laws are cited in the following manner: Title of Act Year, Chapter Number (where the chapter is optional); "Computer Misuse Act 1990, c. 18" is an example.

There is a growing number of legislative and regulatory requirements in the United States and around the world, but there are two overarching U.S. laws that you should be familiar with:

- U.S. Computer Security Act of 1987
- U.S. Federal Information Security Management Act (FISMA) of 2002

U.S. Computer Security Act of 1987

The Computer Security Act was enacted by the U.S. Congress in 1987 with the objective of improving the security and privacy of sensitive information stored on U.S. federal government computers. The act contains provisions that require establishment of minimally acceptable security practices for federal government computer systems, as well as establishment of security policies for government agencies to meet those practices. As part of this act, security awareness training was established as a requirement for any federal government employee using government computer systems.

The Computer Security Act establishes that the National Institute for Standards and Technology, an agency within the U.S. Department of Commerce, is responsible for setting computer security standards for unclassified, nonmilitary government computer systems, while the National Security Agency (NSA) is responsible for setting security guidance for classified government and military systems and applications.

The Computer Security Act of 1987 was repealed by the Federal Information Security Management Act (FISMA) of 2002, which is discussed next.

U.S. Federal Information Security Management Act (FISMA) of 2002

The Federal Information Security Management Act, commonly referred to as FISMA (pronounced "fizz-muh"), is a U.S. law enacted in 2002 that greatly extends the Computer Security Act of 1987. FISMA acknowledges the importance of information security to the United States' economic and national security interests and requires that

all U.S. federal government agencies and nongovernment organizations that provide information services to these agencies conduct risk-based security assessments that align with the NIST Risk Management Framework (RMF).

Industry Standards and Other Compliance Requirements

Aside from national, state, and local laws and regulations, your organization may be required to comply with certain regulations and standards based on your industry or the type of services you provide. The most prominent industry standards that you should be aware of include the following:

- U.S. Sarbanes–Oxley Act of 2002 (SOX)
- System and Organization Controls (SOC)
- Payment Card Industry Data Security Standard (PCI DSS)

U.S. Sarbanes–Oxley Act of 2002

Following several high-profile corporate and accounting scandals, the SOX was enacted in the United States to reestablish public trust in publicly traded companies and public accounting firms. SOX required companies to implement a wide range of controls intended to minimize conflicts of interest, provide investors with appropriate risk information, place civil and criminal penalties on executives for providing false financial disclosures, and provide protections for whistleblowers who report inappropriate actions to regulators.

Under SOX, the Public Company Accounting Oversight Board (PCAOB) was established as a nonprofit organization responsible for overseeing the implementation of SOX. PCAOB's "Auditing Standards" identify the role that information systems play in maintaining financial records and requires auditors to assess the use of IT as it relates to maintaining and preparing financial statements. As part of PCAOB standards, auditors should broadly consider information security risks that could have a material impact on a company's financial statements. Even though SOX is largely a financially focused law, the regulation has a real and growing impact on IT and information security.

System and Organization Controls

Often confused with SOX (discussed previously), SOC stands for System and Organization Controls and is an auditing framework that gives organizations the flexibility to be audited based on their own needs. There are three commonly used types of SOC audits and reports, aptly named SOC 1, SOC 2, and SOC 3. The three audit and report types align with standards outlined in Statement on Standards for Attestation Engagements

(SSAE) 18, which was published by the American Institute of Certified Public Accountants (AICPA) in 2017 (with amendments made via SSAE 20 in 2019).

- **SOC 1:** An audit and compliance report that focuses strictly on a company's financial statements and controls that can impact a customer's financial statements. A company that performs credit card processing is likely to require a SOC 1 audit and compliance report.
- **SOC 2:** An audit and compliance report that evaluates an organization based on AICPA's five "Trust Services principles": privacy, security, availability, processing integrity, and confidentiality. Many organizations undergo SOC 2 auditing and present a SOC 2 report to regulators and customers to demonstrate compliance with industry standard security controls.
- **SOC 3:** This is a "lite" version of a SOC 2 report and abstracts or removes all sensitive details. A SOC 3 report generally indicates whether an organization has demonstrated each of the five Trust Services principles without disclosing specifics (like exactly what they do or don't do). Companies make SOC 3 reports available to the public and restrict SOC 2 reports to trusted parties.

Payment Card Industry Data Security Standard

If your organization handles payment card information (i.e., credit or debit cards), you are likely required to demonstrate PCI DSS compliance. PCI DSS is a proprietary security standard established in 2004. PCI DSS establishes technical and operational requirements for merchants and service providers that accept or process cardholder data and/or sensitive authentication data, as well as for software developers and manufacturers of the applications and devices used in payment card transactions.

NOTE The Payment Card Industry Security Standards Council (PCI SSC) was formed in late 2006 with the goal of ongoing management of the PCI DSS. The PCI SSC is composed of MasterCard Worldwide, Visa International, American Express, Discover Financial Services, and Japan Credit Bureau. To learn more about the Council and the PCI DSS, visit www.pcisecuritystandards.org.

The PCI DSS includes more than 200 security controls organized into 12 requirements, further categorized into 6 goals that generally align with security best practices. Per the PCI SSC, the PCI DSS covers the following:

- Build and Maintain a Secure Network

  Requirement 1: Install and maintain a firewall configuration to protect cardholder data.

Requirement 2: Do not use vendor-supplied defaults for system passwords and other security parameters.

- Protect Cardholder Data

  Requirement 3: Protect stored cardholder data.
  Requirement 4: Encrypt transmission of cardholder data across open, public networks.

- Maintain a Vulnerability Management Program

  Requirement 5: Use and regularly update antivirus software or programs.
  Requirement 6: Develop and maintain secure systems and applications.

- Implement Strong Access Control Measures

  Requirement 7: Restrict access to cardholder data by business need to know.
  Requirement 8: Assign a unique ID to each person with computer access.
  Requirement 9: Restrict physical access to cardholder data.

- Regularly Monitor and Test Networks

  Requirement 10: Track and monitor all access to network resources and cardholder data.
  Requirement 11: Regularly test security systems and processes.

- Maintain an Information Security Policy

  Requirement 12: Maintain a policy that addresses information security for employees and contractors.

Although PCI DSS is not yet a legal requirement, it is often a contractual requirement, and a prime example of an industry standard that is used to mandate, enforce, and audit security standards for applicable organizations across almost all jurisdictions. Because it is not a legislation, PCI DSS is not governed by or enforced by any government body. Instead, compliance with PCI DSS is assessed and enforced by the payment card companies (e.g., Visa, Mastercard, American Express, etc.) mentioned earlier in this section. Failure to satisfy PCI DSS requirements can cost an organization its ability to receive and process such payment card transactions.

Privacy Requirements

Privacy entails limiting access to personal information to authorized parties for authorized uses; in essence, privacy is maintaining the confidentiality of personal information specifically (as opposed to all sensitive data, in general). As more and more of our personal data moves online, privacy has become one of the biggest security-related concerns for regulators, organizations, and users.

Personal information such as your name, address, and Social Security number is considered personally identifiable information, which must be kept confidential. PII is often

subject to some combination of contractual and regulatory privacy requirements. While the source of the requirements may vary, the seriousness with which organizations should take these requirements does not change. As a CISSP, you must know what PII and other personal data your organization handles, and you must understand all legal, contractual, and regulatory requirements that govern the privacy of that data.

UNDERSTAND LEGAL AND REGULATORY ISSUES THAT PERTAIN TO INFORMATION SECURITY IN A HOLISTIC CONTEXT

As a CISSP, you must be aware of the legal and regulatory requirements that pertain to information security — both broadly and within your particular industry and/or geographic regions. Having a strong understanding of legal and regulatory issues involves being familiar with the security threats that face information systems as well as the national, state, and local regulations that govern your organization's handling of sensitive data and systems. For both the CISSP exam and the "real world," you must be familiar with the laws and regulations that govern handling of cybercrimes and data breaches, licensing and intellectual property handling, import/export controls, transborder data flow, and (of course) privacy.

NOTE Misdemeanor and felony are two legal terms that you'll see throughout this section; these two terms describe criminal acts of varying degrees. In U.S. law, a *misdemeanor* is any "lesser" criminal act that is punishable by less than 12 months in prison. Prison time is often (but, not always) substituted with fines, probation, or community service are often (not always) for misdemeanor charges. A *felony*, under U.S. law, is a more serious criminal offense that carries more serious penalties, including jail time over 12 months (and as high as one's lifetime). In other countries, such as France, Germany, and Switzerland, serious offenses (i.e., "felonies" in the United States) are described as crimes, while less serious offenses are called *misdemeanors* or *delicts*. Other countries, such as Brazil, use the term *contravention* to describe less serious offenses.

Cybercrimes and Data Breaches

A *cybercrime* is any criminal activity that directly involves computers or the internet. In a cybercrime, a computer may be the tool used to execute the criminal activity, or it may be the target of the criminal activity. There are three major categories of cybercrimes:

- **Crimes against people:** These crimes include cyberstalking, online harassment, identity theft, and credit card fraud.

- **Crimes against property:** Property in this case may include information stored within a computer, or the computer itself. These crimes include hacking, distribution of computer viruses, computer vandalism, intellectual property (IP) theft, and copyright infringement.
- **Crimes against government:** Any cybercrime committed against a government organization is considered an attack on that nation's sovereignty. This category of cybercrime may include hacking, theft of confidential information, or cyber terrorism. Hacktivism is another cybercrime that involves hackers seeking to make a political statement with their attacks. Hacktivists often target government entities but may also target other organizations with whom they disagree.

A *data breach* is a specific cybercrime where information is accessed or stolen by a cybercriminal without authorization. The target of a data breach is the information system and the data stored within it. Data breaches, and cybercrimes more broadly, may pose a threat to a person, a company, or an entire nation. As such, there are many laws that govern and regulate how cybercrimes are prevented, detected, and handled.

As a CISSP, you should be familiar with the following global cybercrime and information security laws and regulations:

- U.S. Computer Fraud and Abuse Act of 1986
- U.S. Electronic Communications Privacy Act (ECPA) of 1986
- U.S. Economic Espionage Act of 1996
- U.S. Child Pornography Prevention Act of 1996
- U.S. Identity Theft and Assumption Deterrence Act of 1998
- USA PATRIOT Act of 2001
- U.S. Homeland Security Act of 2002
- U.S. Controlling the Assault of Non-Solicited Pornography and Marketing (CAN-SPAM) Act of 2003
- U.S. Intelligence Reform and Terrorism Prevention Act of 2004
- The Council of Europe's Convention on Cybercrime of 2001
- The Computer Misuse Act 1990 (U.K.)
- Information Technology Act of 2000 (India)
- Cybercrime Act 2001 (Australia)

NOTE Many of the regulations in this section have been around for decades. While most of them are still relevant as of this book's writing, the legal landscape is dynamic and changes every year.

U.S. Computer Fraud and Abuse Act of 1986, 18 U.S.C. § 1030

The U.S. Computer Fraud and Abuse Act of 1986 is the oldest and, yet, still possibly the most relevant cybercrime law currently in effect in the United States. The law has been revised over the years, and you should be familiar with both its original form and the revisions discussed in this section.

The Computer Fraud and Abuse Act (CFAA) is a cybercrime bill that was enacted in 1986 as an amendment to the Comprehensive Crime Control Act of 1984. The CFAA was created to clarify definitions of computer fraud and abuse and to extend existing law to include intangible property such as computer data. Although the CFAA now covers all computing devices, the original law was written to cover "federal interest computers" — a term that was changed to "protected computers" in a 1996 amendment to the act. Section 1030(e)(2) defines a protected computer as one that is

- "[E]xclusively for the use of a financial institution or the United States Government, or, in the case of a computer not exclusively for such use, used by or for a financial institution or the United States Government and the conduct constituting the offense affects that use by or for the financial institution or the Government"
- "[U]sed in or affecting interstate or foreign commerce or communication"

In plain English, a *protected computer* is a computer used by the U.S. government or financial institutions, or one used for interstate and foreign communications and financial transactions. It's important to note here that this definition is broad enough to apply to any computer that is "used in or affecting" government and commerce — a computer does not need to be directly used or targeted by a cybercriminal to be considered protected under this definition.

The CFAA establishes seven criminal offenses related to computer fraud and abuse and identifies the penalties for each:

- **Obtaining national security information:** §1030(a)(1) describes the felony act of knowingly accessing a computer without or in excess of authorization, obtaining national security or foreign relations information, and willfully retaining or transmitting that information to an unauthorized party.
- **Accessing a computer and obtaining information:** §1030(a)(2) describes the misdemeanor act of intentionally accessing a computer without or in excess of authorization and obtaining information from a protected computer. This crime

is upgraded to a felony if the act is committed to gain commercial advantage or private financial gain, if the act is committed in furtherance of any other criminal or tortious act, or if the value of the obtained information exceeds $5,000.

- **Trespassing in a government computer:** §1030(a)(3) extends the definition of trespassing to the computing world and describes a misdemeanor act of intentionally accessing a nonpublic protected computer, without authorization, and affecting the use of that computer by or for the U.S. government. §1030(a)(2) applies to many of that same cases that §1030(a)(3) could be charged, but §1030(a)(2) may be charged even when no information is obtained from the computer. In other words, section 1030(a)(3) protects against simply trespassing into a protected computer, with or without information theft.
- **Accessing to defraud and obtain value:** §1030(a)(4) was a key addition to the 1984 act, and it describes the felony act of knowingly accessing a protected computer without or in excess of authorization with the intent to fraud. Under §1030(a)(4), the criminal must obtain anything of value, including use of the information if its value exceeds $5,000. The key factor with §1030(a)(4) is that it allows information theft (described in §1030(a)(2)) to be prosecuted as a felony if there is evidence of fraud.
- **Damaging a computer or information:** §1030(a)(5) was originally written to describe the felony act associated with altering, damaging, or destroying a protected computer or its information, or preventing authorized use of the computer or information, such that it results in an aggregate loss of $1,000 or more during a one-year period. This provision was later rewritten and now more generally describes a misdemeanor act associated with knowingly and intentionally causing damage to a computer or information. §1030(a)(5) upgrades the crime to a felony if the damage results in losses of $5,000 or more during one year, modifies medical care of a person, causes physical injury, threatens public health or safety, damages systems used for administration of justice or national security, or if the damage affects 10 or more protected computers within 1 year.
- **Trafficking in passwords:** §1030(a)(6) establishes a misdemeanor and prohibits a person from intentionally trafficking computer passwords or similar information when such trafficking affects interstate or foreign commerce or permits unauthorized access to computers used by or for the United States. The term *traffic* here means to illegally transfer or obtain control of a password with the intent to transfer it to another party. This definition is important because it excludes mere possession of passwords if there is no intent to transfer them.
- **Threatening to damage a computer:** §1030(a)(7) describes a felony offense associated with the computer variation of old-fashioned extortion. This provision

prohibits threats to damage a protected computer or threats to obtain or reveal confidential information without or in excess of authorization with intent to extort money or anything else of value.

The U.S. Computer Fraud and Abuse Act of 1986 has seen numerous amendments over time, both directly and through other legislations. Minor amendments were made in 1988, 1989, and 1999, with major amendments being issued in 1994, 1996, and 2001 through various other acts discussed later in this chapter.

U.S. Electronic Communications Privacy Act of 1986

The Electronic Communications Privacy Act (ECPA) was enacted by the U.S. Congress in 1986 to extend restrictions on government wire taps to include computer and network-based communications (rather than just telephone calls). The ECPA complements the CFAA by prohibiting eavesdropping, interception, and unauthorized monitoring of all electronic communications (including those sent over computer networks).

The ECPA does, however, make certain exceptions that allow communications providers (like an ISP) to monitor their networks for legitimate business reasons if they first notify their users of the monitoring. This sets up a legal basis for network monitoring, which has been criticized over the years. The USA PATRIOT Act (discussed in a later section) made several extensive amendments to the ECPA in 2001.

U.S. Economic Espionage Act of 1996

The Economic Espionage Act (EEA) was enacted by the U.S. Congress and signed into law by President Clinton in 1996. The EEA was the first federal law to broadly define and establish strict penalties for theft or unauthorized use of trade secrets. The EEA makes it a criminal offense to copy, download, upload, alter, steal, or transfer trade secrets for the benefit of a foreign entity. The EEA establishes penalties for economic espionage that include fines up to $10 million and imprisonment up to 15 years, as well as forfeiture of any property used to commit economic espionage or property obtained as a result of the criminal act.

U.S. Child Pornography Prevention Act of 1996

The Child Pornography Prevention Act (CPPA) was issued in 1996 to restrict and punish the production and distribution of child pornography on the internet.

U.S. Identity Theft and Assumption Deterrence Act of 1998

The Identity Theft and Assumption Deterrence Act was enacted in 1998, and formally established identity theft as a criminal act under U.S. federal law. Under the act, *identity theft* is "knowingly transfer[ring] or us[ing], without lawful authority, a means of

identification of another person with the intent to commit, or to aid or abet, any unlawful activity that constitutes a violation of Federal law, or that constitutes a felony under any applicable State or local law." Prior to the act, identity theft was not regulated or investigated as a crime, which made it difficult to prosecute the growing number of identity theft claims stemming from the rise of the internet.

USA PATRIOT Act of 2001

The Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism (USA PATRIOT) Act, commonly known as the Patriot Act, was signed into law in 2001 in response to the terrorist attacks that took place in the United States on September 11, 2001. The act was initially issued as a temporary measure, but most measures were reauthorized in 2006.

The Patriot Act amends many of the provisions within the CFAA and the ECPA with both new definitions of criminal offenses and new penalties for previously and newly defined computer crimes.

The Patriot Act attempts to strengthen provisions in the CFAA and ECPA to give law enforcement further authority to protect the United States against terrorist acts. The act has been heavily debated since its inception, with some of the act's provisions having been declared unconstitutional by various federal district courts. Of the act's remaining provisions, the following are particularly relevant to the CISSP exam and to you as a security professional:

- **Section 202 — Authority to intercept wire, oral, and electronic communications relating to computer fraud and abuse offenses:** This section amends the CFAA to authorize investigators to obtain a wiretap for felony violations relating to computer fraud and abuse.
- **Section 209 — Seizure of voicemail messages pursuant to warrants:** This section authorizes investigators to seize voicemail messages with a search warrant. Prior to the Patriot Act, voicemail was only authorized for seizure with a harder-to-obtain wiretap order.
- **Section 210 — Scope of subpoenas for records of electronic communications:** This section updates previous law and grants access to additional information when filing a subpoena for electronic records.
- **Section 212 — Emergency disclosure of electronic communications to protect life and limb:** This section grants special provisions to allow a communications provider (like an ISP) to disclose customer information to law enforcement in emergency situations, such as imminent crime or terrorist attack. Prior to this amendment, communications providers may have been subject to civil liability suits for providing such information without the customer's consent.

- **Section 214 — Pen register and trap and trace authority under FISA:** A *pen register* is a device that shows the outgoing calls made from a phone, while a *trap and trace device* shows incoming numbers that called a phone; these capabilities are often consolidated into a single device called a *pen/trap device*. This section of the Patriot Act authorizes use of these devices nationwide (as opposed to an issuing court's jurisdiction) and broadens authority to include computer and internet-based communications.
- **Section 217 — Interception of computer trespasser communications:** This section amends previous law to allow communications providers and other organizations to allow law enforcement to intercept and monitor their systems. Prior to this amendment, companies were authorized to monitor their own systems, but were not permitted to allow law enforcement to assist in such monitoring.
- **Section 220 — Nationwide service of search warrants for electronic evidence:** This section authorizes nationwide jurisdiction for search warrants related to electronic evidence, such as email.
- **Section 808 — Definition of federal crime of terrorism:** The official definition of terrorism includes, among other things, "destruction of communication lines, stations, or systems."
- **Section 814 — Deterrence and prevention of cyberterrorism:** This section strengthens penalties associated with violations in the CFAA, including doubling the maximum prison sentence from 10 to 20 years.
- **Section 815 — Additional defense to civil actions relating to preserving records in response to government requests:** This amendment absolves an organization from civil penalties associated with violations of the ECPA if the organization is responding to "a request of a governmental entity."
- **Section 816 — Development and support for cybersecurity forensic capabilities:** This section requires the U.S. Attorney General to establish regional computer forensic laboratories to support forensic examinations on seized or intercepted computer evidence. Section 816 also requires these laboratories to provide forensic analysis training and education to federal, state, and local law enforcement personnel and prosecutors. This section also includes open-ended language authorizing these forensic labs "to carry out such other activities as the U.S. Attorney General considers appropriate."

U.S. Homeland Security Act of 2002

The Homeland Security Act was enacted in 2002, building off the Patriot Act's response to the September 11, 2001, terrorist attacks in the United States. The

Homeland Security Act sparked the largest U.S. government reorganization since the creation of the Department of Defense in 1947. Under the Homeland Security Act, dozens of government agencies, offices, and services were consolidated into the newly created U.S. Department of Homeland Security (DHS). With the creation of the DHS, a new cabinet-level position, Secretary of Homeland Security, was also created. Title X of the Homeland Security Act identifies several standards, tactics, and controls that should be used to secure U.S. federal government information. Title X and its subsections establish the authorities, responsibilities, and functions associated with information security.

U.S. Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003

The U.S. Controlling the Assault of Non-Solicit Pornography and Marketing Act was signed into law in 2003. This law established the United States' first national standards for sending commercial emails in response to the growing number of complaints over spam (unwanted) emails. The law requires companies to allow email recipients to unsubscribe or opt out from future emails and establishes a variety of requirements around email content and sending behavior. CAN-SPAM designates the Federal Trade Commission (FTC) as responsible for enforcing the provisions within the Act.

U.S. Intelligence Reform and Terrorism Prevention Act of 2004

The Intelligence Reform and Terrorism Prevention Act of 2004 established the National Counterterrorism Center (NCTC) and the position of the Director of National Intelligence (DNI). Under this law, the Department of Homeland Security and other U.S. government agencies are required to share intelligence information to help prevent terrorist acts against the United States. This act also established the Privacy and Civil Liberties Oversight Board with the intent of protecting the privacy and civil liberties of U.S. citizens.

The Council of Europe's Convention on Cybercrime of 2001

The Convention on Cybercrime, also known as the Budapest Convention, is the first international treaty established to address cybercrime. The treaty was first signed in 2001 and became effective in 2004, and has since been signed by more than 65 nations (the United States ratified the treaty in 2006). The treaty aims to increase cooperation among nations and establish more consistent national laws related to preventing and prosecuting cybercrime.

The Computer Misuse Act 1990 (U.K.)

The Computer Misuse Act came into effect in the United Kingdom in 1990 and introduced five offenses related to cybercrime:

- Unauthorized access to computer material
- Unauthorized access with intent to commit or facilitate commission of further offenses
- Unauthorized acts with intent to impair, or with recklessness as to impairing, operation of computer, etc.
- Unauthorized acts causing, or creating risk of, serious damage
- Making, supplying, or obtaining articles for use in other offenses

Information Technology Act of 2000 (India)

The Information Technology Act was passed by the Indian Parliament in 2000 and amended in 2008. The act established legal recognition of electronic documents and digital signatures, while it also established definitions and penalties for cybercrimes such as data theft, identity theft, child pornography, and cyber terrorism.

Cybercrime Act 2001 (Australia)

The Cybercrime Act 2001 was Australia's response to the September 11, 2001, terror attacks in the United States. The Cybercrime Act 2001 defined serious computer offenses such as unauthorized access, unauthorized modification, and unauthorized impairment of electronic communication, and also established penalties for such crimes.

Licensing and Intellectual Property Requirements

Despite the growing list of cybercrime laws that exist today, it's still fairly difficult to legally define and prosecute computer crimes. As a result, many prosecutors fall back on traditional criminal law concepts such as theft and fraud. No matter what organization you work for, there is a good chance that you have some sort of IP that needs to be protected against theft and fraud. IP may include software, data, multimedia content like music and movies, algorithms, drawings, and so much more. As a CISSP, it is your job to protect all forms of IP.

There are various organizations around the world that establish and protect IP rights; among them are the World Trade Organization (WTO), World Customs Organization (WCO), and the World Intellectual Property Organization (WIPO).

There are numerous intellectual property laws and regulations in the United States, and they fit into five categories:

- Licensing
- Patents
- Trademarks
- Copyrights
- Trade secrets

Licensing

Legal protections over intellectual property allow creators and inventors to profit from their work. Unfortunately, the ease with which information can be duplicated and transmitted has made it easier for people to copy information in violation of the legitimate owner's rights.

From an economic perspective, the effect is tremendous. By 2022, the global trade in counterfeited and pirated products, both physical and online, will grow to between 1.9 and 2.8 trillion dollars. Estimates by the Business Software Alliance (BSA) suggest that more than 40 percent of the software in use worldwide is not properly licensed.

Counterfeit goods also present significant economic as well as physical risks. A $460 billion–a–year industry, counterfeiting has been simplified by the e-commerce platforms and expedited international shipping, which has accompanied the lowering of trade barriers. The secondary impacts of illegal use of intellectual property are equally surprising. One estimate suggests that 23 percent of all bandwidth is consumed by activities that infringe on intellectual property.

While emerging technologies present opportunities for improving licensing methods, lack of enforcement remains one of the largest hurdles. With more applications transitioning to a cloud-enabled model, ensuring legal software licensing goes hand in hand with software as a service.

The use of unlicensed software increases the risk of software vulnerabilities, as the users are unable to get patches and updates. This leaves the users of bootleg software at risk when compromises are found in the software. While the vendors patch their legitimate versions, the unlicensed versions don't get the updates. It is somewhat ironic that by illegally using unlicensed software, individuals are more likely to be targeted by other illegal actors. The effect of this was seen most clearly in the rapid distribution of the WannaCry malware in China, where estimates suggest that 70 percent of computer users in China are running unlicensed software, and state media acknowledged that more than 40,000 institutions were affected by the attack.

Patents

A *patent* is a government-issued license or grant of property rights to an inventor that prohibits another party from making, using, importing, or selling the invention for a set period of time. In the United States, patents are issued by the United States Patent and Trademark Office (USPTO) and are usually valid for 15 or 20 years. To qualify for a patent, an invention must be new, useful, and nonobvious. Patents issued by the USPTO are only valid in the United States and its territories; inventors must file patent applications in all countries where they want to be protected under national patent law. There is a European Patent Office (EPO), Eurasian Patent Organization (EAPO), and African Regional Intellectual Property Organization (ARIPO), among others. As a CISSP, you should familiarize yourself with the local IP laws in your jurisdiction.

United States patent law is codified in 35 U.S.C. and 37 C.F.R. and enforced by the U.S. legal system (not the USPTO). For international violations of a U.S. patent, a patent holder may pursue action by the U.S. International Trade Commission (ITC) instead of or in addition to the court system; the ITC can issue exclusion or cease and desist orders to restrict the infringed product from entering the United States. The most robust remedy for international infringement of a U.S. patent may only be achieved through the courts.

Trademarks

According to the USPTO, a *trademark* is "a word, phrase, symbol, and/or design that identifies and distinguishes the source of the goods of one party from those of others." A *service mark* is a similar legal grant that identifies and distinguishes the source of a service rather than goods. The term *trademark* is commonly used to refer to both trademarks and service marks. Think of the brand name Coca-Cola as a popular trademark; the word Coca-Cola distinguishes that specific brand from Pepsi or any other brand of cola/soda/pop. In addition to 35 U.S.C., trademarks are protected under the Trademark Law Treaty Implementation Act (U.S. Public Law 105-330). Unlike patents, a trademark does not expire after a set period of time. Instead, trademark rights last as long as the mark is used in commerce, which can be as long as forever.

Copyrights

A *copyright* is a legal protection granted to the authors of "original works of authorship" that may include books, movies, songs, poetry, artistic creations, and computer software, among other categories. Copyrights created by an individual are protected for the life of the author plus 70 years. Copyright law in the United States was last generally revised by the Copyright Act of 1976 and codified in 17 U.S.C. The U.S. Copyright Office handles registration, recording, and transferring of copyrights, although an original work does not need to be registered to receive copyright protections.

Trade Secrets

A *trade secret* is a proprietary formula, process, practice, or combination of information that a company has exclusive rights to. Using an earlier example, the recipe that Coca-Cola has maintained since 1886 is a trade secret because it is proprietary and has economic value to the company only because of its secrecy. In the United States, trade secret laws are generally left up to the states, although most states have adopted the Uniform Trade Secrets Act (UTSA), which was last amended in 1985. In addition, the Economic Espionage Act of 1996 (discussed earlier in this chapter) and the Defend Trade Secrets Act (DTSA) of 2016 both establish the theft or misappropriation of trade secrets as a federal crime.

Import/Export Controls

Many countries closely regulate the movement of technology through their borders. This might be done to protect local industries from external competition, limit the exportation of sensitive technologies (like encryption), or meet other policy goals of a particular nation. As a CISSP, you should be aware of the implications of any import/export controls in which your organization operates or to which your company's employees may travel.

NOTE The United States, European Union, and other jurisdictions sometimes issue *sanctions* (government edicts that prohibit doing business with a given person, group, organization, or country) against particular countries or particular entities. These sanctions come and go much more frequently than import/export laws and can pose challenges for security teams that operate in or do business with sanctioned entities. As a CISSP, you should be aware of sanctions that impact your organization and help ensure your organization's IT systems meet relevant legal requirements.

One of the most well-known regulations that establishes import/export controls is the U.S. International Traffic in Arms Regulations (ITAR). ITAR regulates the export of defense articles and defense services to keep those sensitive materials out of the hands of foreign nationals. ITAR applies to both government agencies and contractors or sub-contractors who handle regulated materials outlined in the United States Munitions List (USML). Regulated products and technical data include satellites, aircraft, spacecraft, missiles, and much more. Merely sending an email containing ITAR-controlled data (like a blueprint or 3D design file) is considered an export under ITAR. As such, it's important that your organization maintains proper security controls to restrict the flow of ITAR data to legitimate people and locations.

The European Union also places restrictions on dual-use technology. ECPA No. 428/2009 of May 5, 2009, requires member states to participate in the control of exports, transfer, brokering, and transit of dual-use items. In 2017, these regulations were updated to reflect controls over cyber weapons.

A number of countries have adopted laws or regulations that require security reviews to be conducted or, in some cases, denied companies the authority to import products to their countries altogether. In 2016, China passed a broad cybersecurity law that requires information technology vendors to submit their products to the Ministry of State Security for technical analysis. The law allows the ministry to demand source code for inspection as part of the review process. Similar expectations have been placed on software products by Russia and other nations. In 2017, the U.S. government, citing security concerns, singled out Kaspersky Labs, legislating that the company's products would not be allowed on any U.S. government computer system.

Transborder Data Flow

The concept of transborder data flow is closely related to the previously discussed topic of import/export controls. More specifically, this concept focuses on requirements around restricting certain data to or from specific geographic locations or jurisdictions. The ITAR discussed in the previous section is a great example of a legislation that restricts the flow of data. Under ITAR, data must remain within the United States; otherwise, it is considered an export (which may or may not be permitted). Further, ITAR specifically prohibits regulated data from being sent to Iran, Syria, North Korea, and other specified countries. ITAR requirements are particularly noteworthy for public cloud infrastructures that have a global footprint. Many cloud providers have developed the concept of "GovCloud" or similar regionalized cloud offerings to support ITAR and other import/export requirements that restrict transborder data flow.

Many jurisdictions require that certain types of data must be processed inside their borders. This trend has been increasing in recent years, on the assumption that the information, by default, will be more secure, will be available to governments on legal request, and will have the economic benefit of inducing operators of data processing centers to locate facilities within their countries. More than 34 countries have some sort of data localization requirement.

Data localization law took on greater importance following the Snowden disclosures of the range of collection activities performed by the National Security Agency (NSA). Data localization laws were seen as providing some protection against the intelligence activities of foreign powers.

The economic argument for data localization is not necessarily convincing. A substantial body of research suggests that the costs of barriers to data flows in terms of lost

trade and investment opportunities, higher IT costs, reduced competitiveness, and lower economic productivity and GDP growth are significant. The estimates suggest that localization reduces the GDP by 0.7 to 1.7 percent in Brazil, China, the European Union, India, Indonesia, Korea, and Vietnam.

Nevertheless, many countries (in addition to the United States, as already mentioned) have adopted such laws.

Russia

In 2015, Russia became one of the first regimes to require all data collected inside Russia on Russian citizens to be stored inside Russia. The regulations implementing the law may not require localization if the information service is not directed at Russia (i.e., use of Russian language, use of Russian top-level domains, etc.); this has still had significant impact on information providers. Some providers, including Google, Apple, and Twitter, have acquired computing capabilities in Russia to comply with the law. Others, most notably LinkedIn, have resisted the law, and their services have been blocked or curtailed inside Russia.

China

In China, the enforcement of the Cybersecurity Law will place new restrictions on the movement of information. China has asserted sovereignty over the internet operating within its borders and has installed network protections, including limiting access points and strict firewall rules to censor data made available inside China. Article 37 of the Cybersecurity Law requires network operators in critical sectors to store all data that is gathered or produced by the network operator in the country on systems in the country. In particular, the law requires data on Chinese citizens gathered within China to be kept inside China and not transferred abroad without the permission of the Chinese government.

Privacy

Privacy and information security go hand in hand. As discussed earlier in this chapter, privacy is effectively the security principle of confidentiality applied to personal data. There are several important regulations around the globe that establish privacy and data protection requirements. As a security professional, it's important that you understand each privacy regulation that governs your jurisdiction. As a CISSP, you may be familiar with the following regulations, among others, depending on your jurisdiction:

- U.S. Federal Privacy Act of 1974
- U.S. Health Insurance Portability and Accountability Act (HIPAA) of 1996

- U.S. Children's Online Privacy Protection Act (COPPA) of 1998
- U.S. Gramm-Leach-Bliley Act (GLBA) of 1999
- U.S. Health Information Technology for Economic and Clinical Health Act (HITECH) of 2009
- Data Protection Directive (EU)
- Data Protection Act 1998 (UK)
- Safe Harbor
- EU-US Privacy Shield
- General Data Protection Regulation (GDPR) (EU)

NOTE The Asia-Pacific Economic Cooperation (APEC) Privacy Framework is intended to provide member nations and economies with a flexible and consistent approach to information privacy protection without unnecessarily stifling information flow. Although it's not a law or regulation, the APEC Privacy Framework aims to improve information sharing with a common set of privacy principles and is worth reading if you do business in an APEC member economy.

U.S. Federal Privacy Act of 1974, 5 U.S.C. § 552a

The Federal Privacy Act is a U.S. law that was enacted in 1974. The Privacy Act establishes and governs practices related to the collection, maintenance, use, and dissemination of PII by U.S. government agencies. The purpose of the Privacy Act is to balance the government's need to maintain information about citizens and permanent residents with the rights of those individuals to keep their personal information private. Among its provisions, the Privacy Act states that "no agency shall disclose any record which is contained in a system of records by any means of communication to any person, or to another agency, except pursuant to a written request by, or with the prior written consent of, the individual to whom the record pertains." Although the Privacy Act of 1974 substantially predates the internet, the provisions within the act continue to remain relevant and manifest in the form of online privacy consent forms and other mechanisms used to serve as "written consent of the individual."

NOTE Criminal violations of the Federal Privacy Act are deemed misdemeanors and may be subject to penalties of up to $5,000 per violation.

U.S. Health Insurance Portability and Accountability Act of 1996

HIPAA was signed into law in 1996, while the HIPAA Privacy Rule and Security Rule each went into effect in 2003. Organizations that must comply with HIPAA requirements are known as *covered entities* and fit into three categories:

- **Health plans:** This includes health insurance companies, government programs like Medicare, and military and veteran's health programs that pay for healthcare.
- **Healthcare providers:** This includes hospitals, doctors, nursing homes, pharmacies, and other medical providers that transmit health information.
- **Healthcare clearinghouses:** This includes public and private organizations, like billing services, that process or facilitate the processing of nonstandard health information and convert it into standard data types. A healthcare clearinghouse is usually the intermediary between a healthcare provider and a health plan or payer of health services.

The HIPAA Privacy Rule establishes minimum standards for protecting a patient's privacy and regulates the use and disclosure of individuals' health information, referred to as *protected health information*. Under HIPAA, an individual's PHI is permitted to be used strictly for the purposes of performing and billing for healthcare services and must be protected against improper disclosure or use.

The HIPAA Security Rule establishes minimum standards for protecting PHI that is stored or transferred in electronic form. The Security Rule operationalizes the Privacy Rule by establishing the technical, physical, and administrative controls that covered entities must put in place to protect the confidentiality, integrity, and availability of electronically stored PHI (or e-PHI).

Civil penalties for HIPAA violation may include fines that range from $100 to $50,000 per violation, with a maximum penalty of $1.5 million per year for similar violations. Criminal penalties include fines up to $250,000 and potential imprisonment up to 10 years.

U.S. Children's Online Privacy Protection Act of 1998

The Children's Online Privacy Protection Act of 1998 is a U.S. federal law that establishes strict guidelines for online businesses to protect the privacy of children under the age of 13. COPPA applies to any organization around the world that handles the data of children residing in the United States and also applies to children that reside outside of the United States, if the company is U.S.-based. The law sets requirements for seeking parental consent and establishes restrictions on marketing to children under the age of 13.

NOTE According to the Federal Trade Commission (FTC), civil penalties of up to $43,280 may be levied for each violation of COPPA.

U.S. Gramm-Leach-Bliley Act of 1999

The Gramm-Leach-Bliley Act, also known as the Financial Services Modernization Act of 1999, is a U.S. law that requires financial institutions to safeguard their customer's PII. Among the provisions within GLBA, the Financial Privacy Rule requires that financial institutions provide each customer with a written privacy notice that explains what personal information is collected from the customer, how it is used, and how it is protected. The GLBA Safeguards Rule requires organizations to implement proper security controls to protect their customers' personal data.

Penalties for noncompliance with GLBA can include civil fines of up to $100,000 per violation for an organization, and up to $10,000 for officers and directors of a financial services company. In addition, criminal violations of GLBA can include revocation of licenses and up to five years in prison.

U.S. Health Information Technology for Economic and Clinical Health Act of 2009

The Health Information Technology for Economic and Clinical Health Act, referred to as the HITECH Act, was enacted under the American Recovery and Reinvestment Act of 2009. The HITECH Act was created to promote the expanded use of electronic health records (EHRs). Along with increased adoption, the act anticipated an increase in security and privacy risks. As such, the HITECH Act extended HIPAA privacy protections by improving security and privacy protections for healthcare data by imposing tougher penalties for HIPAA compliance violations. Under the HITECH Act, maximum financial penalties were raised to $1.5 million per violation category, per year.

The HITECH Act also introduced a new HIPAA Breach Notification Rule. Under this rule, covered entities are required to disclose a breach of unsecured protected health information to affected parties within 60 days of discovery of the breach. In addition to notifying affected individuals, the Breach Notification Rule requires covered entities to report breaches affecting 500 or more people to the U.S. Department of Health and Human Services and a major media outlet servicing the jurisdiction of the affected parties.

Data Protection Directive (EU)

The Data Protection Directive, officially known as Directive 95/46/EC, was enacted by the European Parliament in 1995. The Data Protection Directive aimed at regulating the

processing of the personal data of European citizens. Although it has since been superseded by the GDPR (discussed in a later section), the Data Protection Directive was the first major privacy law in the European Union and is considered the foundational privacy regulation in all of Europe.

Data Protection Act 1998 (UK)

The Data Protection Act was established by the United Kingdom Parliament to enact the provisions within the EU's Data Protection Directive. The Data Protection Act established that UK citizens held the legal right to control their personal information and was designed to enforce privacy of personal data stored on computing systems. The Data Protection Act 1998 was later superseded by the Data Protection Act 2018, which was designed to enforce and supplement provisions within the GDPR (discussed in a later section).

Safe Harbor

The International Safe Harbor Privacy Principles, often short-handed as just "Safe Harbor," is an agreement between the United States and European Union, established between 1998 and 2000, that was developed to reconcile differences between U.S. and EU privacy laws. Under Safe Harbor, a U.S. company could self-certify that it met data privacy requirements agreed upon by the United States and European Union. Safe Harbor was ruled invalid by the European Court of Justice in 2015 and replaced with the EU-US Privacy Shield soon after.

EU-US Privacy Shield

The EU-US Privacy Shield was the second attempt by the European Union and United States to agree upon principles to mutually regulate the exchange of personal data between the two jurisdictions. The agreement was reached in 2016, less than a year after Safe Harbor was ruled invalid by the European Court of Justice. By 2020, however, the same court declared the EU-US Privacy Shield invalid.

General Data Protection Regulation (EU)

The GDPR is considered by most to be the world's strongest data privacy law. GDPR was established in 2016 and replaced the EU's 1995 Data Protection Directive with hundreds of pages of regulations that require organizations around the world to protect the privacy of EU citizens. With this sweeping regulation, companies around the world that do business with European customers have been forced to rethink their approach to data security and privacy. As a CISSP and information security leader, this is one legislation that you'll likely need to be familiar with.

NOTE If your organization stores or processes the personal data of EU citizens or residents, then GDPR applies to you, whether or not your company is located in the EU.

GDPR Article 5 establishes and describes seven principles for processing personal data:

- **Lawfulness, fairness, and transparency:** Obtain and process personal data in accordance with applicable laws and fully inform the customer of how their data will be used.
- **Purpose limitation:** Identify "specific, explicit, and legitimate" purpose for data collection, and inform them of such purpose.
- **Data minimization:** Collect and process the minimum amount of data necessary to provide the agreed-upon services.
- **Accuracy:** Ensure that personal data remains "accurate and where necessary kept up-to-date."
- **Storage limitation:** Personal data may be stored only long as necessary to provide the agreed-upon services.
- **Integrity and confidentiality:** Ensure appropriate security of personal data, and provide protection against unauthorized access, and accidental loss or destruction. This includes implementing data anonymization techniques to protect your customers' identities, where necessary.
- **Accountability:** The *data controller* (i.e., the party that stores and processes the personal data) must be able to demonstrate compliance with all of these principles. Many customers pursue industry-standard certifications, like ISO 27001, to demonstrate accountability and commitment to security and privacy.

TIP Article 17 within the GDPR establishes a person's "right to be forgotten." This provision grants the *data subject* (i.e., the person whose data is being used) the right to have their personal data deleted if one of several circumstances exists and is a critical concept that information security professionals must consider when developing their data storage and retention policies.

GDPR Chapter 4 contains several articles that establish requirements related to the data controller and processor and requires that *data processors* (i.e., an organization that stores and processes PII on behalf of a data controller) prioritize security and privacy. Of particular interest, Article 25 requires "data protection by design and by default"; this is a huge directive that codifies what security professionals have been recommending as best practice for years.

NOTE GDPR Article 33 establishes rules that require data controllers to notify proper authorities within 72 hours of becoming aware of a personal data breach.

✔ GDPR Fines

The GDPR imposes stiff fines on data controllers and processors for noncompliance.

Determination

Fines are administered by individual member state supervisory authorities (83.1). The following 10 criteria are to be used to determine the amount of the fine on a noncompliant firm:

- **Nature of infringement:** Number of people affected, damage they suffered, duration of infringement, and purpose of processing
- **Intention:** Whether the infringement is intentional or negligent
- **Mitigation:** Actions taken to mitigate damage to data subjects
- **Preventative measures:** How much technical and organizational preparation the firm had previously implemented to prevent noncompliance
- **History:** Past relevant infringements, which may be interpreted to include infringements under the Data Protection Directive and not just the GDPR, and past administrative corrective actions under the GDPR, from warnings to bans on processing and fines
- **Cooperation:** How cooperative the firm has been with the supervisory authority to remedy the infringement
- **Data type:** What types of data the infringement impacts; see special categories of personal data
- **Notification:** Whether the infringement was proactively reported to the supervisory authority by the firm itself or a third party
- **Certification:** Whether the firm had qualified under-approved certifications or adhered to approved codes of conduct
- **Other:** Other aggravating or mitigating factors, including financial impact on the firm from the infringement

Lower Level

Up to €10 million, or 2 percent of the worldwide annual revenue of the prior financial year, whichever is higher, shall be issued for infringements of:

- Controllers and processors under Articles 8, 11, 25–39, 42, 43
- Certification body under Articles 42, 43
- Monitoring body under Article 41(4)

Upper Level

Up to €20 million, or 4 percent of the worldwide annual revenue of the prior financial year, whichever is higher, shall be issued for infringements of:

- The basic principles for processing, including conditions for consent, under Articles 5, 6, 7, and 9
- The data subjects' rights under Articles 12–22
- The transfer of personal data to a recipient in a third country or an international organization under Articles 44–49
- Any obligations pursuant to member state law adopted under Chapter IX
- Any noncompliance with an order by a supervisory authority

UNDERSTAND REQUIREMENTS FOR INVESTIGATION TYPES

In this section, we compare and contrast different investigation types, including administrative, criminal, civil, and regulatory investigations. For each investigation type, we discuss who performs the investigation, the standard for collecting and presenting evidence, and the general differences between the types.

In discussing legal matters, it is important to stress that laws and courts vary significantly across the globe; there are a great many particular distinctions between how law enforcement, courts, lawyers, and judges behave and perform, depending on where you live, where the events leading to the investigation occurred, and other variables. The information presented in this book is largely based on traditions of English common law, strictly as an example; however, it is absolutely essential that you, as a security

professional, familiarize yourself with the laws and regulations relevant to your locale and customers so you can provide adequate, informed service.

Burden of proof is the requirement that the criminal prosecutor or civil plaintiff/claimant prove the claims they are making against the accused, or defendant. The party making a claim must demonstrate the truth of that claim, with compelling evidence; the entity defending against the claim, in most modern societies, is presumed innocent or without fault — that is, the court will not recognize the validity of a claim against anyone until that claim is substantiated and the defendant is proven guilty. The amount and strength of proof required to sway the judgment away from this presumption of innocence differs depending on which kind of claim is being made; for instance, whether the claim is being made by one private party against another or whether the claim is being made by the government against a person or organization (more on this distinction in just a moment). In the U.S. legal system, the two predominant standards of proof that must be met are called preponderance of the evidence and beyond a reasonable doubt.

Preponderance of the evidence is the lower standard of the two and is used primarily in civil actions. It essentially means that the evidence shows that the defendant is more likely to have caused the damage than not. In other words, the evidence convinced the judge, jury, or ruling body that there was at least a 51 percent chance that the defendant caused the damage.

The second standard, *beyond a reasonable doubt*, is much harder to prove and is used primarily in criminal actions. It is insufficient for the evidence to merely make the judge or jury *lean* more toward guilt than not. In this case, the evidence has to be so clear and compelling that a "reasonable" person has no doubt or reservation about the defendant's guilt after seeing it.

Administrative

When discussing investigations, for (ISC)^2 purposes, the term *administrative* will refer to actions constrained to those conducted within a single organization — that is, the organization performs an administrative investigation of itself. Internal investigations are typically performed when the matter involves some violation of organizational policy and does *not* involve any external entities such as law enforcement, investors, third-party suppliers, or attackers.

NOTE To avoid confusion, it is important to distinguish how the term *administrative* is used in a variety of ways to avoid confusion. For (ISC)^2, it means an internal investigation. In the United States, *administrative law* refers to a set of laws made by regulatory bodies (such as the Drug Enforcement Agency, the Food and Drug Administration, and the like). For the purposes of the CISSP Body of Knowledge, an administrative investigation will *only* refer to an internal investigation.

The organization itself can task anyone to perform activities for administrative investigations. This can include staff and employees within the organization (physical and IT security personnel, auditors, management, etc.) or might involve specialized contractors hired by the organization to perform investigative tasks.

The burden of proof for administrative investigations is the lowest of all investigation types. Management can use whatever criteria they choose to believe evidence.

Punitive measures that may result from administrative investigations include employee termination, loss of privilege, reassignment, and so forth. Management might also choose to change the type of investigation as a result of findings made during the administrative investigation; if the administrative investigation reveals that the parties involved engaged in intentional/malicious or criminal activity, management may escalate to civil actions (lawsuits) or filing criminal charges, both of which would require investigatory actions relevant to those situations.

Despite the low burden of proof required for management to act in an administrative investigation, care should still be taken during the process. Occasionally, evidence gathered during an administrative investigation may lead to or be used in a civil or criminal investigation, as stated earlier. If evidence is mishandled during an administrative investigation, it may compromise the ability to use that evidence in later proceedings. If there is any uncertainty about whether an administrative investigation may ultimately escalate, a discussion of this concern with management or in-house or outside counsel is prudent.

Consider this example of an investigation: The IT department contacts the security office to make a report of an employee misusing the organization's internet connection to engage in unauthorized file sharing, in direct violation of the organization's policy. The security office makes the situation known to management; management instructs the IT and security departments to gather information about the user's online activity. Personnel in the IT and security departments work together to gather log data about the user's account and machine, and they present this information to management. Management consults with the legal and human resources departments to evaluate courses of action. Management decides to terminate the employee.

This is strictly an administrative investigation.

Criminal

Criminal investigations involve prosecution under criminal laws. The government, at the federal, state, or local level, prosecutes violations of its laws by imposing fines, imprisonment, or, in some extreme cases, even death for offenders. Criminal investigations are conducted by law enforcement organizations, which can include local, state, federal, or even international agencies. While some CISSPs are in law enforcement positions and conduct criminal investigations themselves, most of us will likely be reporting criminal incidents to law enforcement and helping to collect/provide evidence.

For a law enforcement agency to take part in prosecuting a criminal matter, jurisdiction must first be established. Jurisdiction, as we discussed in the section, "Determine Compliance and Other Requirements," is the legal authority of a governmental body (such as a court or enforcement agency) over a specific matter, often based on geography. With crimes that involve information assets, determining jurisdiction can be complicated and frequently may involve several different government bodies, locales, and laws.

Once jurisdiction has been established, the law enforcement investigator first tries to understand what happened, what damage was done, and what possible range of crimes apply for possible prosecution. In some cases, because of the global nature of IT, a case may be dropped or referred to another law enforcement agency due to a combination of jurisdictional issues, the cost of the investigation versus the scale and impact of the crime, and the likelihood of successful prosecution.

As the investigation progresses, law enforcement begins to understand who the potential suspects might be and what evidence is available, and the investigator must begin to narrow the focus to specific laws and statutes. Many countries, provinces, cities, and other jurisdictions have a variety of laws relating to the misuse and abuse of technology.

Typically, criminal courts have the highest legal standard for determining liability and guilt; this is often referred to as evidence that shows that the accused has caused harm beyond a reasonable doubt. With this standard, the overwhelming majority of evidence must show that the defendant is guilty, leaving the court with no other rational conclusion.

The criminal investigator collects evidence until the elements can be proven or until it is clear that they *cannot* be proven. They use investigative techniques, including digital forensics (covered in Chapter 7, "Security Operations"). The investigators may secure media and devices as necessary for evidence.

When gathering evidence, law enforcers may or may not be required to get a court order, allowing the government to access property, devices, and data that are owned by private entities. These court orders may be in the form of warrants or subpoenas; some must be issued by a judge, while others can be issued by any officer of the court (such as a government-commissioned prosecutor).

When a private organization requests law enforcement involvement with or in response to a suspected incident, that organization may give permission to the government to access the property/devices/data, and there is no need for a court order; the organization owns the property/devices/data and therefore can allow access to anyone it chooses. In criminal matters, the security professional in the employ of an organization requesting law enforcement response should *not* try to investigate without guidance from law enforcement personnel. In other words, if you, as a security professional, suspect a crime has been committed and are going to report this crime to the government, you should suspend investigative activity (beyond containment of immediate damage) until

and unless otherwise instructed by the government/authority and immediately escalate the issue to management. It is possible that investigative actions by untrained personnel unfamiliar with legal procedure and the processes for proper evidence collection and handling can taint the evidence and otherwise impede an investigation.

Further, a security professional in the employ of an organization should not unilaterally make the decision to contact law enforcement. This can be a complex decision and should be made in consultation with management and in-house and/or outside counsel.

Lastly, additional rules apply to security professionals or investigators who are employed by law enforcement and prosecutorial agencies. While a company that owns evidence can simply choose to provide that evidence to law enforcement, stringent rules apply to the collection, handling, and analysis of evidence by law enforcement and the prosecution of employees. Government investigators must be conscious of and understand the legal requirements that apply to them; this will include (among other requirements) whether search warrants are necessary to seize evidence, a stringent adherence to chain of custody procedures, and the analysis of evidence that does not exceed what is legally permitted in a given situation.

Civil

Civil law governs relations and interactions between private entities. The plaintiff in a civil case sues for compensation for a loss or relief from some type of dispute. As information security practitioners, we may be called on to support our clients when they are either plaintiffs or defendants in civil suits. The following are examples of possible civil actions that a security professional, like you, may be involved in:

- **Your organization is the plaintiff:** If someone accesses your production environment without authorization and steals data, causing harm to your organization, your organization might sue the perpetrator for damages (restitution for the harm that was caused). You may be called on to oversee collection of evidence (e.g., logs from penetrated hosts, intrusion detection systems, and network appliances) proving the defendant caused the harm. (Note: This may be in *addition* to criminal action brought by the government against the defendant.)
- **Your organization is the defendant:** If a former employee accuses the organization of creating a hostile work environment, you may have to oversee collection of evidence (such as emails between managers and executives discussing how employees are treated), as well as preventing the destruction of potential evidence (referred to as *destruction hold notice*, *preservation notice*, *litigation hold*, or similar terms) upon request by courts or attorneys.

Unlike criminal law, in civil proceedings, the usual standard of proof is preponderance of the evidence, meaning it is a much lower burden of proof. Preponderance of the evidence is a simple majority of fault/liability; if the plaintiff can prove to the court that the defendant is even 50.1 percent culpable for the damages, the defendant will lose the civil case.

In a civil proceeding, there is no question of guilty versus not guilty but rather liable versus not liable. If the defendant is found liable, they may be ordered to pay for damages, to stop an activity that is harming the plaintiff, or to honor a contract or agreement into which they had previously entered. Unlike criminal sentences, a litigant cannot be jailed or put to death for liability in a civil lawsuit. However, if a civil litigant refuses to obey a court order, it can result in a contempt of court charge, which could eventually lead to jail time.

Because the burden of evidence and stakes involved in losing a civil case are much lower than they are in criminal cases, the level of effort in collecting and processing the evidence is likewise lower. This is *not* to say that evidence in civil cases can be handled in a haphazard or careless manner; due care must still be taken to perform actions in a suitable, professional way. However, in civil cases, investigation and evidence collection will not be performed by badged law enforcement personnel and government agents; instead, it is done by information technology and security professionals, such as CISSPs.

Similar to criminal trials, there are rules as to what evidence may be used in a civil trial. Collected evidence that is deemed unreliable may be excluded by a judge presiding over the trial. Care should be taken to retain original copies of evidence collected by an investigator, and chains of custody should be well documented. Original evidence should never be altered, with very few exceptions, and without direct instructions from counsel who is overseeing an investigation or handling the case. Spoliation of evidence (i.e., altering or destruction of the original) can lead to exclusion of evidence in a case or, in some situations, can lead to a separate lawsuit for the damages resulting from the spoliation.

If there is uncertainty about the rules surrounding the collection and handling of evidence for a civil lawsuit, consultation with a digital forensic expert or counsel can be helpful.

Regulatory

Regulatory investigations involve determining whether an organization is compliant with a given regulation or legal requirement. Regulations have the force of law; consequently, regulatory investigations are similar to criminal investigations. Regulations are written under the auspices of protecting the average citizen or consumer, protecting the environment, or making an industry safer and more equitable.

NOTE It is important to understand the (ISC)² definition of regulations, as it is used in the CISSP CBK. A regulation is not a standard, guideline, or suggestion — it is law, established by a government body. For instance, in the United States, the Environmental Protection Agency (EPA) is part of the federal government; the EPA writes regulations concerning activities that may impact the environment (such as handling hazardous/toxic waste, transportation of certain materials, and so forth). EPA regulations have the force of law: anyone violating these regulations may be prosecuted by the government. Conversely, the PCI-DSS is *not* a regulation, as defined by (ISC)²; the PCI-DSS is a contractual standard, affecting only those parties that voluntarily choose to comply with it (i.e., merchants that accept credit card payment).

Government agencies perform regulatory investigations to determine whether sufficient evidence exists to prove some violation of rules or regulations. These agencies have the authority and discretion to decide when to perform investigations. These agencies have their own internal investigators, prosecutors, and courts for their proceedings. Regulators can also demand information from target organizations or utilize audit report data in addition to or instead of performing their own investigations.

The burden of proof for regulatory investigations is the preponderance of the evidence, and the penalties typically involve fines and injunctions. There are, however, instances where regulators call for referral to criminal law enforcement that may result in prison time.

Industry Standards

Investigation is a broad term. There are currently many standards and guidelines offered in this realm, some of which are dependent on the jurisdiction or industry in which the organization operates. This section takes a quick look at some of the most common standards and guidelines that are related to investigations.

ISO/IEC 27043:2015 recommends procedural steps for conducting security incident investigations. These guidelines cover many incident scenarios from the preparation phase all the way through to the conclusion of the investigation. The scenarios covered include incidents such as data loss or corruption, unauthorized access, and confirmed data breaches.

ISO/IEC 27037:2012 provides guidelines for handling digital evidence. This is covered through a four-step process of identification, collection, acquisition, and preservation. Evidence collection and handling is covered across many types of media and scenarios, including magnetic and optical storage media, mobile devices, camera systems, standard computers, and collecting network traffic data from network devices. This publication also covers chain of custody procedures and how to properly exchange evidence between jurisdictions.

NIST SP 800-86, "Guide to Integrating Forensic Techniques into Incident Response," overlaps significantly in terms of content with the two previous sources. It is the NIST perspective on the digital forensic process. It details how to build a forensic capability within your organization, what that means, and which tools and training your staff will need. The publication also describes how to structure forensic policies, standards, and procedures for your organization and what they should contain. Most importantly, NIST SP 800-86 describes the digital forensic process overall in four phases: collection, examination, analysis, and reporting.

NIST SP 800-101 Revision 1, "Guidelines on Mobile Device Forensics," has a self-explanatory title. It covers the unique requirements for acquiring, preserving, examining, analyzing, and reporting on the digital evidence present on mobile devices. The technical differences associated with mobile devices are discussed, such as differences in memory type and file structure that affect evidence collection. The publication also discusses sensitive areas that may arise when a mobile device is privately owned.

As you can see from this list of industry standards, evidence management and digital forensics are at the heart of conducting technology-based investigations. Domain 7 of the CISSP CBK covers Security Operations, and we further discuss understanding and complying with investigations in Chapter 7 of this book.

DEVELOP, DOCUMENT, AND IMPLEMENT SECURITY POLICY, STANDARDS, PROCEDURES, AND GUIDELINES

Although technical security controls like firewalls, encryption, and sophisticated access control mechanisms are incredibly important in maintaining the security of your organization's data, documents such as policies, standards, procedures, and guidelines are the most essential components of an information security program. Each of these documents is different, yet they are closely related and work together to guide your organization's behavior. Figure 1.3 shows the relationship that policies, standards, procedures, and guidelines have with each other.

Policies

A *policy* is a formal set of statements that establish a system of principles to guide decisions and actions. More specifically, a *security policy* is a set of statements that identifies the principles and rules that govern an organization's protection of information systems and data. Policies can be company-wide, system-specific, or issue-specific (e.g., an incident response policy). Some common examples of security policies include the following:

- Acceptable use policy
- Access control policy
- Change management policy
- Remote access policy
- Disaster recover policy

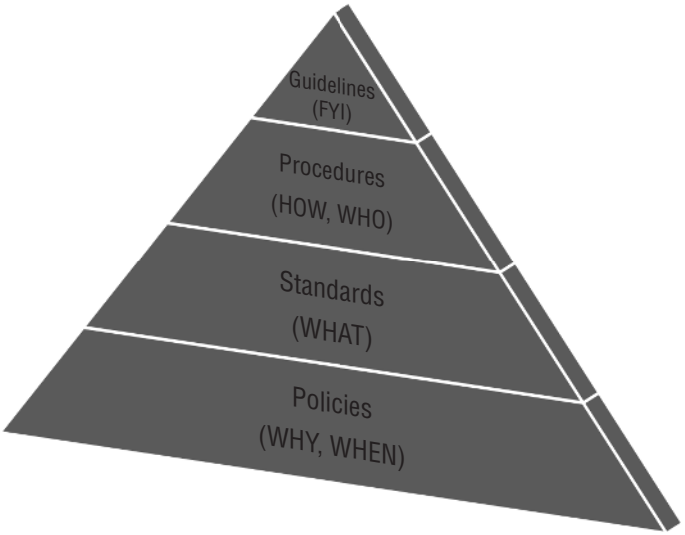

FIGURE 1.3 **Relationship between policies, procedures, standards, and guidelines**

Policies set the foundation for your organization's security program and are typically written to be broad enough to be applicable and relevant for many years. Much like the foundation of a building, security policies should survive long-term and are less likely to change than other documents, although they should be periodically reviewed and updated, as necessary. Standards, procedures, and guidelines are supporting elements that provide specific details to a complement an organization's policies.

Standards

Standards are specific and granular requirements that give direction to support broader, higher-level policies. Standards establish specific behavior and actions that must be followed and enforced to satisfy policies. Standards may be mandatory for a given organization, if mandated by contract or law. The Federal Information Processing Standards (FIPS), for example, are publicly announced standards that were developed by NIST to establish various security requirements for U.S. government agencies; FIPS 140-2, for example, is a standard that establishes security requirements for cryptographic modules.

Within this standard, there are specific encryption devices that are permitted and prohibited for use within U.S. government systems.

Baselines are related to standards and establish a minimum level of a security for a system, network, or device. For example, your organization might maintain individual baselines for each operating system that the company uses. Each of these baselines should identify what specific settings, applications, and configurations must be in place to meet your company's security standards and policies. While not specifically called out in the CISSP CBK, you should be familiar with baselines and understand how they fit into the bigger picture.

As a subset of baselines, *security baselines* express the minimum set of security controls necessary to safeguard the CIA and other security properties for a particular configuration. Scoping guidance is often published as part of a baseline, defining the range of deviation from the baseline that is acceptable for a particular baseline. Once scoping guidance has been established, then tailoring is performed to apply a particular set of controls to achieve the baseline within the scoping guidance. Scoping and tailoring is further discussed in Chapter 2, "Asset Security."

Procedures

A *procedure* is a detailed step-by-step guide to achieve a particular goal or requirement. Procedures tell you *how* to implement your policies and *how* to meet your standards and baselines. Some common examples of security procedures include the following:

- Vulnerability scanning procedures
- Backup and restore procedures
- Account provisioning procedures
- Patch management procedures

As a CISSP, you may be called upon to create, update, and manage information security policies at your organization. In addition, as a CISSP, you must ensure that other, noninformation security procedures (e.g., HR and transaction processing procedures) within your organization are safe, secure, and compliant with relevant policies and standards.

Guidelines

A *guideline* is similar to a standard but is a recommendation rather than a mandatory requirement. Guidelines refer to policy statements and offer flexible suggestions for meeting the intent of the policy or recommendations to implement the requirements in standards and baselines.

There are many sources of guidelines for information security practice. Certainly, the CISSP CBK is one, as it reflects a broad range of security practices but is not prescriptive inside an organization's information security environment. The ISO/NIST/ITIL frameworks are often leveraged as guidelines; however, they may become policies or standards if the organization has a compliance expectation. Other sources of guidelines include manufacturers' default configurations, industry-specific guidelines, or independent organizations such as the Open Web Application Security Project (OWASP) work in software development.

IDENTIFY, ANALYZE, AND PRIORITIZE BUSINESS CONTINUITY REQUIREMENTS

Business continuity (BC) and disaster recovery (DR) (discussed in detail in Chapter 7) are closely related concepts that help an organization continue essential operations during a security incident and recovery from a disaster (or major disruptive event) as quickly and securely as possible. Business continuity and disaster recovery are quite often referred to in the same breath, but it's important that you understand the role that each plays.

- A *business continuity plan* (BCP) is a methodology and set of protocols that deals with allowing an organization to keep their key business functions running in the event of a crisis; this is sometimes referred to as *continuity of operations* (COOP). Business continuity includes all of the preventative controls and the management of employees that help preserve the functionality of the overall business during a disaster.
- A *disaster recovery plan* (DRP) is the set of processes that deal with restoring your information systems and operations, securely and efficiently, after a disruptive event occurs. DR is the subset of BC whose primary objective is to minimize business downtime and reclaim normal operations as soon as possible.

NOTE Generally speaking, BCP is broadly focused on all critical business functions and operations, while disaster recovery is more narrowly focused on systems, applications, and data. For example, BCP covers everything from DDoS and ransomware attacks (discussed in Chapter 3) to natural disasters that shut down entire datacenters. DRP, on the other hand, focuses on getting things back to "normal" — "things" here includes both IT systems and business processes.

When a disaster occurs, BC and DR activities are each kicked off simultaneously — business continuity tasks keep essential business functions running while disaster recovery

actions work toward getting things back to normal. For both BC and DR planning, a business impact analysis (BIA) can help identify critical business functions.

Business Impact Analysis

According to the ISO, a *business impact analysis* is "the process of analyzing the impact over time of a disruption on the organization." In other words, a BIA helps an organization identify its essential business functions and understand the impact that a disaster would have on each of those functions; the BIA provides the primary justification for the business continuity plan and its requirements. The BIA helps an organization identify which of its business functions are more resilient and which are more fragile.

To complete the BIA, you should begin by establishing your BC project team, scope, and budget; we cover this step in the next section, "Develop and Document the Scope and the Plan." You should ensure that you have executive support for BCP activities. BCP does not yield an immediate or tangible return, so you need senior leadership's support for financial resources, staffing, and overall strategic vision.

The next step is one of the most important: identify all your critical business functions (CBFs) and other essential business elements. The key word here is *business*, as you should be focused on identifying the essential functions that are critical to your business operations, whether your business involves selling widgets or saving lives.

Identifying CBFs requires input from a broad range of stakeholders. The perspectives of the system owners, subject-matter experts, customers, and suppliers all help in identifying an organization's potential CBFs. A list of CBFs and essential business elements should include the following:

- Personnel
- Business processes
- Information systems and applications
- Other assets

For each CBF that your organization identifies, you should perform a risk analysis to identify any vulnerabilities that exist, along with steps to mitigate those weaknesses. In addition, you must determine the likelihood of adverse events affecting your CBFs and the level of impact the business is willing to accept due to disruption to any one of the critical business functions. Determining the level of impact of a disaster is done in several ways, and there are a few metrics that you should be comfortable with:

- *Maximum tolerable downtime* (MTD), or maximum acceptable outage (MAO), expresses the total length of time a critical business function can be unavailable

without causing significant, long-term harm to the business; it is the longest time a CBF can remain disabled before it threatens the organization's long-term survival. MTD must be defined by the system owner, who is ultimately responsible to the organization for the proper operation of the CBF. Exceeding the MTD is an expression of unacceptable risk by the business owner.

- *Recovery time objective* (RTO) is the planned time necessary to restore a system to the point where it meets the minimum service expectations of the system owner. In other words, RTO is the maximum period of time within which a CBF must be restored after a disruption to avoid unacceptable business consequences. Since unacceptable disaster occurs when the MTD is exceeded, the RTO, by definition, must be less than or equal to the MTD. The RTO must be adjusted by the application of additional controls to bring it within the MTD. At the point where the business owner is no longer willing to apply control, they have accepted the risk of operation.
- *Recovery Point Objective* (RPO) represents the measurement of tolerable data loss, represented as a period of time. As with the MTD, this must be defined by the business, and the business is responsible for resourcing the controls to achieve the RPO.

MTD, RTO, and RPO are essential thresholds for business continuity planning. Figure 1.4 visually represents these concepts and how they fit together.

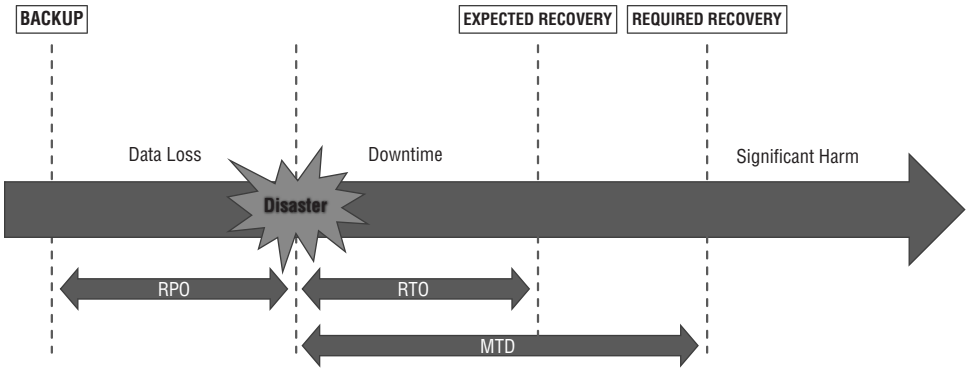

FIGURE 1.4 **Relationship between MTD, RTO, and RPO**

NOTE MTD, RTO, and RPO are important planning horizons and estimates that are strongly linked to an organization's risk tolerance. During an actual recovery event, organizational leaders can and usually do adapt in real time, making these thresholds more guidelines than strict survival limits.

Develop and Document the Scope and the Plan

The BCP itself is the organization's commitment to maintaining the operations of the business, and the steps the organization takes to do so. This plan focuses on the people, processes, and technologies on which the business relies to deliver goods and services to its customers. The information derived from your BIA activities should be used to document the scope of your business continuity plan.

The BCP must protect an organization's critical business functions and its customers and provide the capability for an organization to continue effective business operations at a service level and in a time period that meets any legal and regulatory requirements in addition to the organization's defined MTD, RTO, and RPO (discussed in the previous section).

The scope of the BCP must encompass all of the organization's operations, including each business area and within every geographic region that the organization does business. While there is no one-size-fits-all for business continuity planning, the scope of most plans includes the following:

- Critical business functions
- Threats, vulnerabilities, and risks
- Data backup and recovery plan
- BCP personnel
- Communications plan
- BCP testing requirements

Once your organization has completed a business impact assessment, you should have a list of CBFs and an understanding of your organization's threshold for downtime and loss for each of them. The next phase of continuity planning involves identifying the specific mechanisms and procedures to mitigate risks to your CBFs and maintain compliance with your established MTD, RTO, and RPO.

As with any good plan, a BCP involves people, processes, and technologies — in that order. In next three sections, we cover some of the requirements and techniques involved in protecting these three categories of assets.

People

People are always, without exception, your most valuable and critical asset. The first goal of any BCP must be to ensure the safety of your people during and after an emergency. In the context of BCP, "people" include your employees, contractors, customers, vendors, and any other living human being that may be affected by an adverse event.

After ensuring the safety of your people, you must ensure that they are provided with the resources necessary to continue working as normally as possible. This may include shelter (e.g., an alternate work site) and food they require to survive and complete their BCP and operational tasks.

A well-designed business continuity plan must include protocols for notifying all affected people (internal and external) that an adverse event has occurred. You must ensure that multiple methods of communications are in place to notify critical BCP personnel, in case one or more methods are unavailable due to the disaster. Further, management and key BCP stakeholders must receive regular status updates during a disaster to provide awareness and allow strategic decisions to be well-informed.

Processes

The BCP team must evaluate every critical business function and determine what resources must be available during a disaster. Your continuity plan should identify the critical supplies and logistics required to maintain critical operations, and it should establish a process to ensure those resources remain continuously available.

One of the most essential BCP processes assures an organization that its critical data processing facilities and capabilities remain operational during a disaster. Your organization should identify where and how you will continue your critical data processing functions. The most relevant method of addressing this topic is by developing processes for the use of alternate sites during a disaster. The primary recover site types are hot sites, cold sites, and warm sites. We cover these in Chapter 7.

Technologies

Hardware and software failures — that's just part of the reality of technology. A business continuity plan must anticipate these failures and outline controls and procedures to mitigate the risk of technology failure. System and data backups are the most tried-and-true way that organizations address this risk. You must have a comprehensive backup process to ensure your critical systems and data are captured, stored, and available for recovery when necessary. You should maintain multiple copies of your most critical information. If your organization maintains on-premise systems (e.g., if you run a data center), for, one set of your backups must be stored offsite; this serves to protect at least one replica of your data in case a disaster destroys your primary location. If you use cloud-based systems, you should maintain backup copies in multiple *cloud regions* (or geographic locations where datacenters are located) so that your data is recoverable at any given time.

Aside from information backup, your BCP should establish a protocol for maintaining redundant systems to continue supporting your business during a significant negative

event. Redundant electrical supplies, water supplies, telecommunication systems, and network connectivity systems are required to ensure continued operations during a disaster. Many organizations lease two or more internet service providers (ISPs), multiple utility providers, and at least two banking providers, in case the disaster originates with one of these providers rather than internal to the organization.

CONTRIBUTE TO AND ENFORCE PERSONNEL SECURITY POLICIES AND PROCEDURES

The Security and Risk Management domain of the CISSP CBK covers many of the foundational concepts necessary to build and manage secure systems and data. Because hardware, software, and technical controls tend to get all the attention, it's important that you keep in mind that the human element is perhaps the biggest part of information security. An essential part of your organization's security planning should be focused on policies and procedures to ensure the security of your employees. In this section, we cover topics such as candidate screening and hiring, employee onboarding and offboarding, managing external personnel (i.e., vendors, consultants, and contractors), and other important personnel security considerations.

Candidate Screening and Hiring

Candidate screening and hiring the right employees is a critical part of assuring the security of your company's systems and data. Not only do you need to make sure to hire the right fit for the job, but it's also critical that you are familiar with a candidate's background and history before bringing them into your organization and giving them access to your sensitive information.

There are a couple things your organization must do before beginning to recruit candidates for a position. First, the hiring manager should work with HR to clearly and concisely document the job description and responsibilities. Having a job description with well-documented responsibilities can help you recruit the right person for the job and can later be used as a measuring stick to assess the employee against the expectations set before they were hired. Next, you should identify the classification or sensitivity of the role, based on the level of damage that could done by a person in that role who intentionally or negligently violates security protocols. The classification or sensitivity assigned to a role (referred to as a *risk designation* by NIST, for example) should inform the types of authorizations an employee will receive once they are hired; as such, the thoroughness of your candidate screening process should match the security of the position that you're

filling. As a CISSP, risk designation (or the equivalent in your jurisdiction) should be considered prior to granting any employee access to sensitive information.

Once a potential employee or contractor is identified, your organization should verify the information in their application and confirm their suitability for the position by conducting a background check. Generally speaking, an employment background check may include the following checks and verifications:

- Education
- Work history
- Citizenship
- Criminal record
- Credit and financial history
- References

In addition to the previous list, candidate screening may include drug testing and/or further investigation for highly sensitive roles, or positions requiring a special security clearance (this is especially relevant for employment with a government agency). As a CISSP, you should ensure that your organization has policies and procedures in place to screen and hire candidates in accordance with any relevant regulations in your jurisdiction.

NOTE While background investigations used to be strictly handled by organizations specifically created to conduct them, many employers have added online background screening to their standard procedures. In these circumstances, an employer may choose to research a potential candidate's social media and online presence to gain a fuller picture of that person's attitude, intelligence, professionalism, and general character. Organizations should have clear policies that define the appropriate uses of internet and social media research, standardize which information is to be taken from the social media sites, verify the accuracy of the information, and disclose to applicants the potential use of internet and social media in deciding which applicants to consider.

Employment Agreements and Policies

When joining an organization, an employee generally signs an employment contract that may include one or more employee agreements that make certain stipulations by which the employee must abide. The most common employee agreements are nondisclosure agreements and noncompete agreements.

A *nondisclosure agreement* (NDA) is an agreement that restricts an employee or contractor (or anyone else with access to sensitive information) from disclosing sensitive information they obtain through the course of their employment or relationship with an organization. An NDA is designed to protect the confidentiality of the organization's data (such as trade secrets or customer information) and is often a lifetime agreement (even after the employee leaves the company).

A *noncompete agreement* is an agreement that restricts an employee from directly competing with the organization during their employment and, in most cases, for a fixed time after employment. Noncompetes are one-way agreements that are designed to protect organizations from unfair competition by former employees or contractors. As an example, if you are hired as a hardware engineer for a mobile phone designer, you may be required to sign a noncompete stating that you will not work for other companies that design mobile phones for at least 18 months after termination of your employment; the idea here is that your inside knowledge of the company will present less of a disadvantage after those 18 months.

In addition to NDAs and noncompete agreements, employees may be responsible for reviewing and/or signing various employment policies such as acceptable use policies, code of conduct, or conflict of interest policies.

Onboarding, Transfers, and Termination Processes

Onboarding, transfers, and termination are three stages of employment that each comes with its own security considerations. The processes that bring people into an organization set the tone for their work behavior. Similarly, employee termination processes should clarify people's obligation to respect the protection of the organization's intellectual property and data security as they leave the company. As a security professional, you should be actively engaged with the business to ensure that onboarding, transfer, and termination processes are clearly documented and set behavior expectations during all stages of employment.

Onboarding

Setting good expectations for work behavior should start before the employee walks in the door. Part of the employee orientation program should address information security expectations and requirements. Employees should be reminded of their obligations to protect information and current threats to the organization's information assets, particularly if they are likely to be the targets of malicious actors. Further, orientation practices should inform new employees of the processes for reporting security incidents, their role in maintaining the security of their work area, and the company's classification and

categorization processes so they can identify the level of control necessary for particular information.

Employees should also be made generally aware of the existence of controls that monitor their use of the organization's assets. Not only does this provide them with assurance that the organization does indeed take action to protect its information, but the information alone may act as a deterrent to inappropriate behavior. The intent is not to provide the employee with sufficient technical detail to defeat the controls, but to make sure they understand that their actions may be scrutinized.

Transfers

Organizations should have well-defined policies and procedures for handling an employee transferring from one role to another. Part of this process should involve reviewing the employee's existing access to information and evaluating the need for continued access to the same information. Where possible, your organization should seek to remove access that will no longer be needed in the employee's new role; this enforces the principle least privilege, which we discussed earlier in this chapter. In addition, you should have a process in place to identify any role-based training that the employee needs to take prior to the transfer; this is particularly critical when the employee's new role comes with new responsibilities or access to information at a higher sensitivity.

Termination

Taking appropriate care when people depart an organization is just as important as ensuring they are properly brought into the organization. Terminations may be voluntary (i.e., an employee retires or finds a new job) or involuntary (i.e., an employee is fired, furloughed, or otherwise "let go"). These former insiders represent a risk to the organization, and appropriate actions must be taken to ensure they do not compromise the operations, intellectual property, or sensitive information with which they have been entrusted.

When an individual leaves an organization on good terms, it is relatively easy to go through the standard checklist: suspending electronic access, recovering their access badges and equipment, accounting for their keys, and changing the key codes on cipher locks that the departing employee used are among many other standard practices. Most organizations have well-structured off-boarding processes to ensure the removal of access when an individual is no longer entitled to organizational information or resources.

Involuntary termination of employment is an emotionally charged event for all involved. In virtually all cases, an involuntary termination forces the employer to assume the terminated individual is a threat to the organization, and appropriate action should be taken to protect organizational assets. Termination procedures at most organizations

include specific processes to notify the information security organization to disable access to electronic and physical systems.

Where possible, recovery of property that an involuntarily terminated employee used should be attempted. Where appropriate, the recovered material should be tracked as evidence and retained for subsequent forensic analysis. Finally, once the individual has left the organization, remaining staff should be informed that the terminated individual is no longer allowed access and that any attempts by that individual to access resources or property should be reported.

TIP It is not unusual for individuals to have taken steps to harm the organization in the event that they were terminated. The most obvious forms of this are the theft of data by the terminated individual, who hopes to sell back the key to the organization (i.e., ransomware), use the information to begin or join a competing organization, or disclose the information to discredit the organization. Strong data security practices and a well-developed insider threat program are essential in defeating malicious activities by terminated employees. User and Entity Behavior Analytics (UEBA), for example, can help detect a disgruntled employee who is heading toward a rage quit.

Vendor, Consultant, and Contractor Agreements and Controls

Many organizations require expertise or talent that does not exist inside their organizations. These relationships may exist for goods or services, but both types of acquisition open the organization to risk. Information security policies should be in place to ensure that these relationships do not expose the organization's sensitive information to an unreasonable amount of risk. NDAs and other employment agreement policies play a big part in establishing expectations with third parties and can lead to additional compliance burden on the organization who must enforce them.

Compliance Policy Requirements

Responsibilities for compliance with applicable policies and regulations should be clearly documented and understood by all employees within an organization. In many cases, employees may be required to sign an attestation stating that they have reviewed and agree to comply with all company policies and applicable regulations.

Employees and other parties with access to systems and information must undergo initial and periodic training that includes security awareness and job-based training. Generally, annual recertification is a good way to ensure that all parties with access remain in compliance with employment policies.

Privacy Policy Requirements

Your organization's *privacy policy* is an explanation of your company's personal data collection and use practices. Privacy policies should link back to applicable privacy laws and regulations, such as HIPAA in the United States and GDPR (discussed later in this chapter) for companies that handle EU residents' information.

Your privacy policy should explain what kind of personal data is collected, how your organization will or won't use it, and how the personal data will be stored, maintained, and secured. The privacy policy should be made available to all personnel, and many organizations require a signed acknowledgment from each employee.

UNDERSTAND AND APPLY RISK MANAGEMENT CONCEPTS

The topic of risk management, and all the concepts within it, is at the heart of information security and is the core of every strong information security program. *Risk management* includes all the processes associated with identifying threats and vulnerabilities and quantifying and addressing the risk associated with those threats and vulnerabilities. Risk management processes provide a structured method for making security decisions such as purchasing and implementing security tools and hiring people. This section covers the key concepts behind risk management and guides you through applying these concepts in your organization.

Identify Threats and Vulnerabilities

In security, a *risk* is the potential for negative impact on the organization, its goals or objectives, or its assets (including people, systems, and data) due to a threat exploiting a vulnerability. You should note that there are dozens of definitions for each of these terms (i.e., risk, threat, and vulnerability) across different industries. We'll discuss these terms further, but it's important to understand that risk lies at the intersection of the three components shown in Figure 1.5.

NOTE There are two classifications of risk that you should be familiar with: inherent risk and residual risk. Simply put, *inherent risk* is the risk present *before* any controls are applied, while *residual risk* is the level of risk that remains *after* controls are in place. The concept of security controls is discussed later in this chapter.

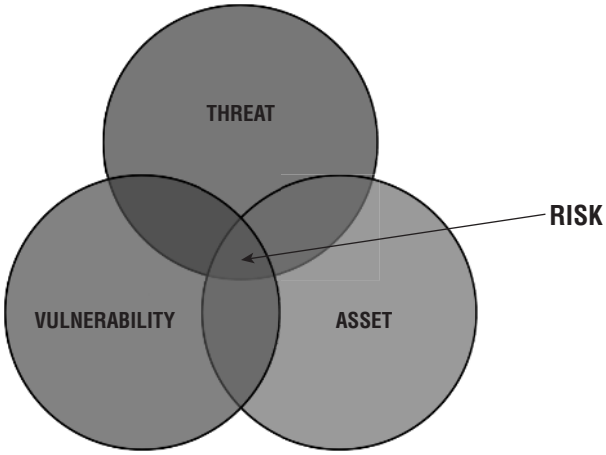

FIGURE 1.5 **Relationship between threats, vulnerabilities, assets, and risks**

Threats

A *threat* is a negative event that can lead to an undesired outcome, such as damage to, or loss of, an asset. A threat is posed by a *threat actor*, which is a person or entity that is capable of intentionally or accidentally compromising an asset's security. As security professionals, threats are the people and events that we work to protect our information and systems from.

Examples of security threats include the following:

- A hacker who wants to encrypt your data and charge you a ransom for it
- A disgruntled employee who wants to steal or sell corporate information
- A fire or other natural disaster that may damage or destroy your datacenter

Vulnerabilities

A *vulnerability* is a weakness or gap that exists within a system that may be exploited (by a threat actor) to compromise an asset's security or trigger a risk event. Vulnerabilities are the things within our systems that we try to fortify and improve.

Examples of security vulnerabilities include the following:

- Unpatched software applications
- Weak access control mechanisms (e.g., weak passwords)
- Faulty fire suppression systems

Assets

An *asset* is anything of value, which may include people, property, and information. Assets are the things that we, as security professionals, are trying to protect. People assets may include your company's employees, contractors, and vendors, as well as your customers. Property assets include tangible things like servers and equipment, as well as intangible things like software code and other intellectual property.

Risk Assessment

Remember that risks are the intersection between threats, vulnerabilities, and assets, as shown in Figure 1.5. A *risk assessment* is the set of activities that involve identifying the threats and vulnerabilities that exist and determining the impact and likelihood of those threats exploiting the identified vulnerabilities.

There are numerous risk frameworks (as discussed in the "Risk Frameworks" section) that provide guidance on conducting risk assessments, but generally speaking, risk assessments include the steps shown in Figure 1.6.

FIGURE 1.6 **Steps for assessing risk**

NOTE The NIST CSF and other modern risk frameworks are recognizing the need for the small to medium business (SMB) community to start with the first risk they identify and manage it, rather than going through the stepwise cycle in Figure 1.6. It's important that you consider your organization's resources and identify a risk management process that works for you.

Risk Identification

The first step in a typical risk assessment process is to identify your assets and determine the value of those assets; this includes identifying and classifying your sensitive data, based on its sensitivity or value to your organization. During the risk identification phase, you find the systems, applications, and information that need protecting and then identify and describe the vulnerabilities and threats that pose a risk to each of those assets.

Risk Analysis

Risk analysis should always begin with a vulnerability assessment (discussed in Chapter 6, "Security Assessment and Testing") and a threat analysis (discussed in the section

"Understand and Apply Threat Modeling Concepts and Methodologies" later in this chapter). This stage of risk assessment is focused on evaluating the likelihood of identified threats exploiting weaknesses (i.e., vulnerabilities) in your environment and determining the impact to your assets if that happens. *Likelihood* describes the probability that an event will occur, and *impact* defines how disastrous the event would be if it were to happen.

Likelihood can be identified by evaluating each threat and assessing the probability that the threats might actually exploit a vulnerability, or weakness. For example, you might determine that the risk associated with a destructive fire is relatively low if you have redundant fire suppression systems that are tested monthly; if you have mechanisms in place to detect and extinguish fires and if you are testing those mechanisms regularly, then the likelihood, or probability, of a fire destroying everything is reduced. Similarly, you might identify insider threat as high likelihood if you've contracted with a large organization without conducting thorough background checks — in this situation, there is a greater probability that something bad will happen.

Impact can be identified by establishing the value associated with each potentially affected asset and determining how that value will be destroyed or otherwise affected by an adverse event. An asset's value can be both *quantitative* (i.e., determined by its cost or market value) or *qualitative* (i.e., determined by its relative importance to you or your organization). By establishing an asset's value, you can better determine the impact of that asset's security being compromised — this allows informed decision-making when determining how much to spend on safeguarding a given resource, as you never want to spend more protecting an asset than the asset itself is worth.

✔ Quantitative Risk Calculation

Risk analysis can be either qualitative or quantitative (or a combination of the two). *Qualitative risk analysis* avoids the use of numbers and tends to be more subjective. *Quantitative risk analysis* is far more precise and objective, because it uses verifiable data to analyze the impact and likelihood of each risk. Quantitative risk calculation involves making measurements to mathematically determine probability (likelihood) and impact. Qualitative risk analysis involves assigning less precise values (like critical, high, medium, and low) to likelihood and impact.

While some risks can be hard to quantify, keep in mind that qualitative analysis can often be vague, imprecise, and even misleading. For example, pandemics were a pretty "low" probability of occurrence prior to 2019, but COVID-19 demonstrated that the overall risk associated with pandemics could be very high.

One important concept in quantitative risk analysis is *annualized loss expectancy* (ALE), which is a metric that helps quantify the impact of a realized threat on your organization's assets. ALE is measured in dollars and is the product of single loss expectancy (SLE) and annual rate of occurrence (ARO), which are each discussed here:

- SLE is a measure of the monetary loss (calculated in dollars) you would expect from a single adverse event. In other words, SLE estimates how much you would lose from one occurrence of a particular realized threat. SLE is calculated by multiplying an asset's value (AV) by its exposure factor (EF). EF is the estimated percentage of loss to a specific asset if a specific threat is realized.
- ARO is the estimated annual frequency of occurrence for a given adverse event. In other words, ARO is the number of times that you expect a particular risk event to occur every year.

Here are the two formulas to keep in mind:

$$ALE = SLE \times ARO$$

$$SLE = AV \times EF$$

Risk Evaluation

During risk evaluation, you compare the results of your risk analysis to your organization's established risk profile or risk tolerance (i.e., how much risk your organization is willing to take on). In doing so, you are able to determine the best course of action for each of your identified risks. We cover the various options for risk response in the following section.

Risk Response/Treatment

Once you identify and assess your organization's threats, vulnerabilities, and risks, you must determine the best way to address each risk; this is known as *risk treatment* (or *risk response*). There are four main categories of risk treatment, as we describe in the following sections: avoid, mitigate, transfer, and accept. Each of these are ultimately leadership/management decisions that should have CISSP input and awareness.

Avoid

Risk avoidance involves eliminating an identified risk by stopping or removing the activity or technology that causes the risk in the first place. Organizations use risk avoidance

when a particular risk exceeds their acceptable risk tolerance, but complete avoidance is often difficult to achieve without business disruption. While this type of risk treatment can often mean simply not doing something, policies that ban the use of removable media or personal cloud storage services are avoidance steps that require upfront investment and action.

Mitigate

Risk mitigation (sometimes called *risk reduction* or *risk modification*) is a strategy that involves reducing the likelihood of a threat being realized or lessening the impact that the realized threat would have on the organization. Risk mitigation is the most common treatment option for identified risks and involves implementing policies and technologies to reduce the harm that a risk might cause. Moving from single-factor to mutifactor authentication is an example of a mitigation treatment for sensitive data access.

Transfer

Risk transference (also known as *risk assignment*) involves shifting the responsibility and potential loss associated with a risk onto a third party. Insurance is the most common form of risk transference. For example, if a company loses customer data due to a cyber breach, the company may rely on their cyber insurance to cover any monetary losses associated with the breach. In this case, the breached organization has transferred financial risk to their cyber insurer, but the company still must manage through some level of reputational risk. It's hard to completely transfer *all* risk, so many people instead use the term *risk sharing*. Using cloud-based services or managed security services is a great example, because risk is split between you, as the customer, and the third-party provider.

Accept

Risk acceptance unsurprisingly involves accepting the risk associated with a particular threat. Risk acceptance is the way to go if avoiding, mitigating, or transferring the risk would cost more than the expected losses of the realized threat. In theory, a risk should be accepted only if it is completely within an organization's risk tolerance. In practice, organizations are often forced to accept potentially painful risks associated with normal business operations.

Countermeasure Selection and Implementation

Mitigation is the most common risk treatment method of the four treatment approaches in the previous section. Risk mitigation involves the selection and implementation of one or more countermeasures (or "security controls") with the goal of reducing the likelihood

of an adverse event or the impact of that event occurring. Countermeasures generally fall into three categories:

- **Personnel-related:** As people are commonly considered to be an organization's "weakest link," these countermeasures often prove invaluable. Hiring (or firing), organization restructuring, and awareness training are some common personnel-related countermeasures. Despite our potential as weaknesses, people in high-performing organizations with strong security awareness programs can often prove to be the greatest security asset.
- **Process-related:** Policy, procedure, and other "workflow-based" mitigations generally fall into this category. As an example, consider the implementation of separation of duties on invoice approval and payment as a process-related mitigation against cyber fraud.
- **Technology-related:** This is the category that typically gets the most attention. Encryption, modifying configuration settings, and other hardware or software changes are common examples of technology-related countermeasures.

When selecting countermeasures, you must consider factors such as security-effectiveness, cost-effectiveness, and operational impact.

Security-Effectiveness

Measuring the security-effectiveness of a security control is an essential step in the selection and implementation process. When selecting your countermeasures, you want to be certain that the specific policy, technology, or operational control that you select is able to directly address a risk identified during your risk analysis process. To do this, one must consider what kind of security risks one wants to prevent, detect, or correct, and then identify countermeasures that specifically target those risks. For example, many security teams choose to throw encryption at everything, but if you are concerned with risks that encryption cannot fix (like availability risks), you are better off using those resources for other countermeasures (such as backups).

Cost-Effectiveness

Perhaps even more important than security-effectiveness (believe it or not), cost-effectiveness is a primary consideration for security teams and the management teams that oversee them. Cost-effectiveness can be calculated by performing a cost-benefit analysis that compares the cost of a countermeasure (or multiple countermeasures) to the costs that would be realized by a compromise of the risks that the countermeasures are intended to mitigate.

A countermeasure can be considered cost-effective if the annual loss expectancy (ALE) *with* the countermeasure plus the cost of countermeasure is less than ALE *without*

the countermeasure. For example, if the ALE associated with theft of sensitive data is $500,000, you can theoretically spend up to $499,999.99 on countermeasures to reduce the ALE of such data theft to $0.01. Of course, you'd want to gain more than a single penny from all your troubles, but this demonstrates the point. Another way to look at it is if the ALE due to ransomware attacks on your company is projected at $200,000 and you spend $50,000 on a sophisticated backup system, the selected countermeasure has a value of $150,000 to your organization, which is quite clearly cost-effective.

NOTE Countermeasures generally have an initial acquisition and implementation cost, followed by recurring (e.g., annual) operating and maintenance costs. You should consider both sets of costs when determining whether a countermeasure makes financial sense for your organization.

Operational Impact

Beyond cost-effectiveness and pure security-effectiveness, you must be sure to evaluate the potential operational impact that a countermeasure may have on your organization. If a countermeasure is too difficult to implement or use, it may have a counterintuitive effect and actually increase risk because it is not being used properly (or at all). For example, some organizations require the use of third-party email encryption platforms to send sensitive information, and some of these platforms are not user friendly at all. Without careful selection of a platform and proper user training, some users may circumvent this countermeasure and send sensitive emails in the clear. Understanding your organization's culture and strategy is an important part of selecting countermeasures that don't have a negative operational impact.

Applicable Types of Controls

A *security control* is any safeguard that is put in place to positively impact security. Security controls may be automatic or manual, and they can be technical (i.e., implemented and executed through hardware, software, or firmware), operational (i.e., related to day-to-day operations and tangible things like security guards, gates, etc.), or management (i.e., implemented by people and related to administrative methods — things like policies, procedures, and guidelines). There are five major types of controls, and you'll notice that some countermeasures (like security guards) may fit into multiple categories:

- **Preventative:** These are the first-line controls that are designed to keep adverse security events from occurring. For example, software applications typically have some form of "input validation" to avoid invalid inputs from being executed and

causing an issue. Firewalls, system backups, and security awareness training are other common examples of preventative controls.

- **Detective:** These controls are designed to identify a negative security event while it is in progress or soon after it occurs. Much like a human detective, this type of control is intended to gather information and help security teams determine what happened, how bad the damage is, and what caused it to happen. Security audits, door alarms, and IDSs are common examples of detective controls.
- **Corrective:** These controls are designed to minimize and repair damages following an adverse security event; they are typically put in place after a detective control identifies a problem. Corrective controls include things such as software patches, configuration file modifications, and new policies that target the cause of the incident.
- **Recovery:** These countermeasures are designed to complement corrective controls, with the intent to get a system back to normal as quickly as possible. Examples include system and data backups and disaster recovery sites.
- **Deterrent:** These controls are designed to discourage attackers by making them think twice about their malicious intents. Wired fences, security guards, and guard dogs are some examples of deterrents.

TIP You should also be familiar with the concept of a *compensating control*, which is a safeguard used in addition to or in place of a primary control; compensating controls are often implemented if a primary control cannot be fully implemented for some reason. For example, if a technical security control is too expensive, you may opt for policies that *encourage* rather than *enforce* a desired behavior. The compensating control may not fully mitigate the risk, but it provides some level of security that wouldn't exist without any control being implemented. PCI-DSS provides some good examples of compensating controls usage.

Control Assessments

Periodic assessment of your security controls is equally as important as the selection and implementation of those controls. In many cases, your organization may have legal or regulatory requirements that dictate how and when to conduct security control assessments (SCA), but in all cases, you should routinely conduct control assessments to ensure that your security and privacy controls remain effective.

SCAs may take the form of self-assessments or external assessments conducted by third parties. There are many different SCA methodologies, but they generally include some form of the following assessment methods: examine, interview, and test. NIST 800-53A, "Assessing Security and Privacy Controls in Federal Information Systems and Organizations," lays out some helpful guidelines for conducting controls assessments and describes the three assessment methods as follows:

- **Examine:** This method is "the process of reviewing, inspecting, observing, studying, or analyzing one or more assessment objects (i.e., specifications, mechanisms, or activities). The purpose of the examine method is to facilitate assessor understanding, achieve clarification, or obtain evidence." Assessors often begin an SCA by requesting a list of artifacts or evidence (such as security policies, configuration files, etc.) that they can examine to form an initial perspective.
- **Interview:** This method is "the process of holding discussions with individuals or groups of individuals within an organization to once again, facilitate assessor understanding, achieve clarification, or obtain evidence." After reviewing any evidence provided during the examine phase, assessors meet with key stakeholders to gain additional clarity on what security controls are in place and how they work.
- **Test:** This method is "the process of exercising one or more assessment objects (i.e., activities or mechanisms) under specified conditions to compare actual with expected behavior." In this stage, an auditor or assessor is seeking to confirm that security controls are implemented as they are documented and that they are operating effectively and as intended.

Chapter 6 covers security assessment extensively.

Monitoring and Measurement

Monitoring and measurement of your controls is an important part of operating a risk-based security program. In addition to conducting periodic (e.g., annual or quarterly) security and privacy control assessments, you should actively and intentionally monitor your controls to measure their effectiveness and assess the health of your overall security program. Depending on your organization's needs, you should develop a set of key performance indicators (KPIs) that allow you to quantify and measure the long-term performance of your controls.

Reporting

Conducting SCAs and other monitoring and measurement activities is useless without a well-managed reporting function. Auditors and assessors generally create formal reports that detail their findings for each control that is assessed. In addition, your security team

should have a process to document and report any important discoveries or metrics to senior leadership, regulators, and other stakeholders.

Some laws, regulations, and industry requirements come with specific reporting guidelines; as an information security leader, you must be familiar with any such requirements that are relevant to your organization. In general, a well-managed risk-based security program includes some level of reporting for the following:

- Internal audits (e.g., self-assessments)
- External audits (i.e., regulator or any other third-party audits)
- Significant changes to the organization's risk posture
- Significant changes to security or privacy controls
- Suspected or confirmed security breaches (or other incidents)

Continuous Improvement

A common goal among security leaders is to continuously improve their organization's security posture and measure their journey toward their desired end state. As a CISSP, you need to continuously identify whether your organization is improving its management of information security risks. You should also seek to continuously improve the return on investment (ROI) associated with the security tools, controls, and processes that your organization implements. There is a fine line between "not secure enough" and "perhaps too many security tools and processes." As a CISSP, you should seek to continuously improve the efficiency of your organization's security management program.

Risk maturity modeling is a process that allows an organization to assess the strength of its security program and create a plan for continuous improvement based on their results. By identifying the maturity of its program on a predefined scale, an organization may better focus on what types of behaviors are necessary to improve, rather than getting caught up strictly in individual security gaps. Maturity models are discussed further in Chapter 8.

Risk Frameworks

A *risk framework* is a structured process for identifying, assessing, and managing an organization's risks. A number of frameworks have been developed to identify and evaluate risk. These frameworks have evolved to address the unique needs of different industries and regulations. Individually, these frameworks address assessment, control, monitoring, and audit of information systems in different ways, but all strive to provide internal controls to bring risk to an acceptable level. While there are several internationally accepted risk frameworks, a number of industry-specific frameworks have also been developed to meet specific needs.

Regardless of the framework, to effectively address risk in an organization, standard processes to evaluate the risks of operation of information systems must take into account the changing threat environment, the potential and actual vulnerabilities of systems, the likelihood that the risk will occur, and the impact to the organization should that risk become realized.

From a governance perspective, the selection of a framework should create a controls environment that is as follows:

- **Consistent:** A governance program must be consistent in how information security and privacy are approached and applied.
- **Measurable:** The governance program must provide a way to determine progress and set goals. Most control frameworks contain an assessment standard or procedure to determine compliance and, in some cases, risk as well.
- **Standardized:** As with measurable, a controls framework should rely on standardization so results from one organization or part of an organization can be compared in a meaningful way to results from another organization.
- **Comprehensive:** The selected framework should cover the minimum legal and regulatory requirements of an organization and be extensible to accommodate additional organization-specific requirements.
- **Modular:** A modular framework is more likely to withstand the changes of an organization, as only the controls or requirements needing modification are reviewed and updated.

There are dozens of different risk management frameworks. While many of the frameworks address specific industry or organizational requirements, you should be aware of the broad characteristics of the more common frameworks.

International Standards Organization

The International Standards Organization has developed the ISO 31000 series of standards to identify principles for general risk management and to provide a set of guidelines for implementation. Developed using the consistent language contained in ISO/IEC Guide 73:2009, the ISO 31000:2018 is intended to be applicable to any organization, regardless of the governance structure or industry. The standard encourages the integration of risk management activities across organizational lines and levels to provide the organization with a consistent approach to management of operational and strategic risks.

ISO 31000:2018 is based on a set of eight principles that drive the development of the risk framework shown in Figure 1.7. That framework, in turn, structures the processes for implementing risk management.

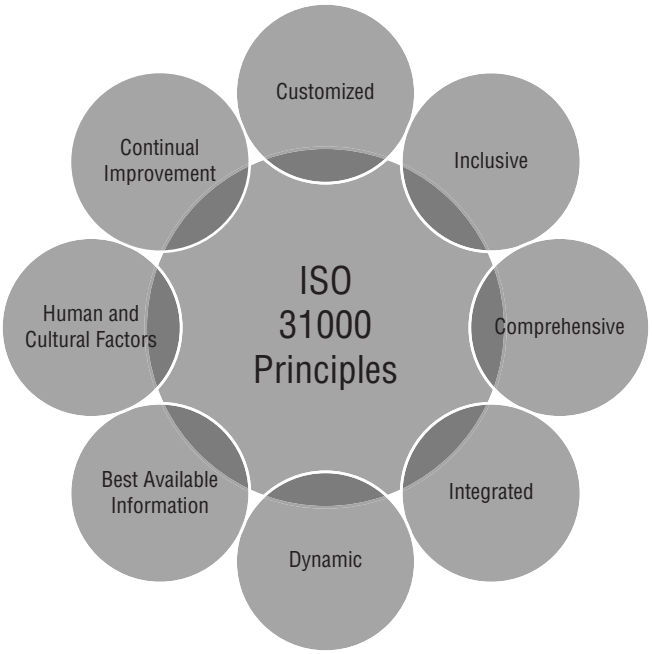

FIGURE 1.7 **ISO 31000:2018**

The eight ISO 31000 principles are described here:

- **Customized:** The framework should be customized and proportionate to the organization and the level of risk.
- **Inclusive:** The appropriate and timely involvement of stakeholders is necessary.
- **Comprehensive:** A structured and comprehensive approach is required.
- **Integrated:** Risk management is an integral part of all organizational activities.
- **Dynamic:** Risk management anticipates, detects, acknowledges, and responds to changes in a timely fashion.
- **Best available information:** Risk management explicitly considers any limitations of available information.
- **Human and cultural factors:** Human and cultural factors influence all aspects of risk management.
- **Continual improvement:** Risk management is continually improved through learning and experience.

To assist organizations in implementing the ISO 31000 standard, ISO 31004, "Risk Management — Guidance for the implementation of ISO 31000," was published to provide a structured approach to transition their existing risk management practices

to be consistent with ISO 31000 and consistent with the individual characteristics and demands of the organization.

While the 31000 series addresses general risk, information security practices are addressed in the ISO 27000 series. The use of the ISO/IEC Guide 73 allows for a common language, but ISO/IEC 27005:2011, "Information technology— Security techniques — Information security risk management," gives detail and structure to the information security risks by defining the context for information security risk decision-making. This context includes definition of the organization's risk tolerance, compliance expectations, and the preferred approaches for assessment and treatment of risk.

ISO 27005 does not directly provide a risk assessment process. Rather, it provides inputs to, and gets outputs from, the risk assessment practice used by the organization. In this framework, the assessment process may be performed in a quantitative or qualitative manner but must be done consistently so that prioritization can be performed. ISO 27005 further emphasizes the need for communication with stakeholders and for processes that continuously monitor for changes in the risk environment.

The ISO standards have seen broad adoption, in part because of the broad international process in the development of the standards. Further, the standards themselves, while constantly under review, connect to other standards managed within the ISO. This enables organizations to adopt those standards that are appropriate for their businesses and provides a more holistic view of an organizations' risk and compliance activities.

U.S. National Institute of Standards and Technology

Through a hierarchy of publications, the National Institute of Standards and Technology provides direction to U.S. government agencies in implementing information security practices. In the current incarnation, the Risk Management Framework (RMF) provides a structured analytical process to identify, control, evaluate, and improve the organization's information security controls. Documented in NIST Special Publication 800-37, "Guide for Applying the Risk Management Framework to Federal Information Systems," it prescribes a six-step process through which the federal government manages the risks of its information systems; the six steps are pictured in Figure 1.8. Though the steps in this framework are tailored to government agencies, they are widely applicable within just about every industry.

The first step of the NIST RMF involves categorizing all information systems based on the potential impact to the organization due to the loss of confidentiality, integrity, or availability. Implied in this process is that the organization must have a comprehensive inventory of systems to apply the categorization standard. Once security categorization has been performed, a baseline set of controls must be selected based on the identified categorization and impact.

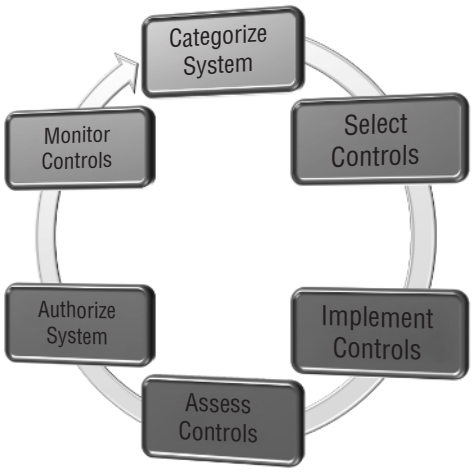

FIGURE 1.8 **NIST Risk Management Framework**

Once the system has been categorized and baseline controls are selected, the controls must be implemented and monitored to ensure that they "are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system." This will produce a set of documents certifying the technical application of the controls.

After categorizing information systems, selecting and implementing controls, and assessing the effectiveness of those controls, organizational leadership then makes a formal decision whether to authorize the use of the system. This decision is based on the ability of the controls to operate the system within the organization's risk tolerance. Finally, the organization must continuously monitor the effectiveness of the controls over time to ensure that the ongoing operation of the system occurs within the organization's risk tolerance.

While focused on the computing activities of the U.S. government, the NIST standards and guidelines have had a pervasive effect on the security community because of their broad scope, their availability in the public domain, and the inclusion of industry, academic, and other standards organizations in the development of the standards. Further, the NIST standards often set the expectations for security practice that are placed on other regulated industries. This is most clearly shown in HIPAA legislation, where healthcare organizations must demonstrate that their controls align with the NIST security practice. Due to its broad reference, the NIST RMF is an important part of the CISSP CBK.

TIP NIST 800-30, "Guide for Conducting Risk Assessments," and the NIST Cybersecurity Framework (discussed in the "Security Control Frameworks" section) both provide practical guidance to help the CISSP frame, present, and inform management decisions about risk.

COBIT and RiskIT

In the late 1990s, the audit community in the United States and Canada recognized that there was a significant gap between IT governance and the larger organizational management structures. Consequently, IT activities were often misaligned with corporate goals, and risks were not comprehensively addressed by the control structure or consistently reflected in financial reporting. To address this gap, ISACA developed a framework through which the IT activities of an organization could be assessed.

The Control Objectives for Information and Related Technology framework differentiates processes into either Governance of Enterprise IT (five processes) or Management of Enterprise IT (32 processes). Each process has a set of objectives, inputs, key activities, and outputs, and measures to evaluate performance against the objectives. As the framework is closely aligned with other management frameworks and tools (ISO 20000, ISO 27001, ITIL, Prince 2, SOX, and TOGAF), it has gained wide acceptance as an encompassing framework for managing the delivery of IT.

Based on the ISACA COBIT governance framework, the RiskIT framework provides a structure for the identification, evaluation, and monitoring of information technology risk. This simplifies the integration of IT risk into the larger organization enterprise risk management (ERM) activities.

The RiskIT framework consists of three domains — risk governance, risk evaluation, and risk response — each of which has three processes. The framework then details the key activities within each process and identifies organizational responsibilities, information flows between processes, and process performance management activities. Additional detail on how to implement the framework and link it to other organizational management practices is contained in the RiskIT Practitioner Guide.

UNDERSTAND AND APPLY THREAT MODELING CONCEPTS AND METHODOLOGIES

Threat modeling is a technique by which you can identify potential threats to your systems and applications, as well as identify suitable countermeasures against those threats.

Threats may be related to overall system vulnerabilities or an absence of necessary security controls. Threat modeling is most often used during the application development phase, but you can also use threat modeling to help reduce risk in existing applications and environments.

The *attack surface* is the total range of areas where an attacker can potentially execute a compromise. With an information system, this might include the methods of communication, the access controls, or weaknesses in the underlying architectures. With a physical environment, the attack surface might include the construction techniques, the location, or the means of entrance and egress. Limiting the attack surface to the minimum number of areas of exposure reduces the opportunities for a threat to become a successful attack.

Threat Modeling Concepts

There are three general approaches to threat modeling. Different risk methodologies, compliance frameworks, or systems prefer one or another, but as a CISSP, you should be able to apply all three approaches to a particular environment. These approaches are attacker-centric, asset-centric, and software-centric (or system-centric), which are each covered in the following sections.

Attacker-centric

The attacker-centric threat modeling approach starts by identifying the various actors who could potentially cause harm to a system. With an attacker-centric approach, you start by profiling a potential attacker's characteristics, skillset, and motivation, and then use that profile to identify attackers who would be most likely to execute specific types of attacks. This approach can be helpful when narrowly approaching a problem by limiting the number of scenarios under analysis. Tactical military intelligence is typically driven by an attacker-centric threat model, as are many business continuity/disaster recovery planning processes. If you work in financial services, you may be familiar with attacker-centric modelling from anti-money laundering (AML) and other anti-financial crimes applications. AML processes involve using process models of how money launderers operate when they attack to identify steps to take in order to thwart such attacks.

Asset-centric

As opposed to an attacker-centric approach, an asset-centric threat model identifies the assets of value first. Assets should be characterized by their value to the organization as well as their value to potential attackers. The means by which the asset is managed, manipulated, used, and stored are then evaluated to identify how an attacker

might compromise the asset. Many compliance regimes focus on protection of an asset (e.g., PHI under HIPAA, PII under the GDPR, or cardholder data under PCI-DSS), so this approach is helpful when establishing or verifying compliance. You'll also find this approach particularly useful when protecting other high-value assets such as intellectual property and security credentials.

Software-centric (or System-centric)

For many information systems environments, the software- or system-centric model is most useful. In this approach, the system is represented as a set of interconnected processes, using architecture diagrams such as dataflow diagrams (DFDs) or component diagrams. These diagrams are then evaluated by threat analysts to identify potential attacks against each component and to determine whether a security control is necessary, exists, and achieves the control effect.

Threat Modeling Methodologies

There are many different threat modeling methodologies. Some of the most important methodologies are STRIDE, PASTA, NIST 800-154, and DREAD — we discuss each of these in the following sections.

STRIDE

STRIDE is a threat modeling methodology developed by Microsoft in the late 1990s to help identify and classify computer security threats. The name itself is a mnemonic for six categories of security threats, discussed here:

- **Spoofing:** *Spoofing* is an attack during which a malicious party assumes the identity of another party (either a user or a system) by falsifying information. A common example of identity spoofing occurs when email spammers modify the "From:" field to depict the name of a sender that the target recipient is more likely to trust. Within applications, spoofing can occur if an attacker steals and uses a victim's authentication information (like username and password) to impersonate them within the application.
- **Tampering:** *Data tampering* is an attack on the integrity of data by intentionally and maliciously manipulating data. Tampering can include altering data on disk, in memory, over the network, or elsewhere. Applications that don't properly validate user input may allow malicious actors to modify values and have the manipulated data stored and used by the application.
- **Repudiation:** *Repudiation* is the ability of a party to deny that they are responsible for performing an action. The threat of repudiation occurs when a user

claims that they did not perform an action, and no other party is able to prove otherwise. In the physical world, signing for a mail delivery is a common form of nonrepudiation — the delivery company maintains a record that you received and accepted the mail on a specific date. In the digital world, an example of a repudiation threat is a user claiming that they did not make an online purchase — even if they did, in fact, make that purchase. Comprehensive logging, digital signatures, and multifactor authentication can be integrated into applications to provide nonrepudiation for high-risk actions.

- **Information disclosure:** *Information disclosure* is when information is shared with an unauthorized party — such as during a data breach or when inadvertently sending an email to the wrong person. This threat compromises the confidentiality of data and carries a great deal of risk depending on the sensitivity of the leaked data. Organizations that store and process PII, PHI, cardholder data, or other confidential information should focus on this threat, and identify controls to mitigate against it. Data encryption, strong access control, and other data protection mechanisms are the keys to protecting against unauthorized information disclosure.
- **Denial of service:** A *denial-of-service* (DoS) attack is a common availability attack that denies access to resources by legitimate users. Controls should be put in place to monitor and detect abnormally high resource consumption by any single user; this may be an indication of either malicious or unintentional resource exhaustion. As a principle, applications should be developed with availability and reliability in mind.
- **Elevation of privilege:** *Elevation of privilege* (or *privilege escalation*) occurs when an unprivileged user is able to upgrade their privileges to those of a privileged user (e.g., a system administrator). Elevation of privilege can give an untrusted party the "keys to the kingdom" and grant them access to and control over sensitive data and systems. Strong access control is required to help protect against this threat. Systems should revalidate a user's identity and credentials prior to granting privileged access, and multifactor authentication should be used, wherever possible.

PASTA

The *Process for Attack Simulation and Threat Analysis* (PASTA) is a risk-based threat model, developed in 2012, that supports dynamic threat analysis. The PASTA methodology integrates business objectives with technical requirements, making the output more easily understood by upper management.

There are seven stages of the PASTA methodology:

- Define objectives
- Define technical scope
- Application decomposition
- Threat analysis
- Vulnerability analysis
- Attack enumeration
- Risk and impact analysis

NIST 800-154

NIST 800-154, "Guide to Data-Centric System Threat Modeling," was released in draft form in 2016. It explicitly rejects that best-practice approaches are sufficient to protect sensitive information, as best practice is too general and often overlooks controls specifically tailored to meet the protection of the sensitive asset. NIST 800-154 establishes four major steps for data-centric system threat modeling:

1. Identify and characterize the system and data of interest.
2. Identify and select the attack vectors to be included in the model.
3. Characterize the security controls for mitigating the attack vectors.
4. Analyze the threat model.

DREAD

DREAD is an older threat modeling technique, previously used by Microsoft but later abandoned. DREAD provides a mnemonic for quantitative risk rating security threats using five categories:

- **D**amage
- **R**eproducibility
- **E**xploitability
- **A**ffected users
- **D**iscoverability

Though it is sparsely used today, you should be familiar with the DREAD mnemonic and the categories that it represents.

Other Models

Other threat modeling methodologies include the following:

- Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) is an approach for managing information security risks, developed at the Software Engineering Institute (SEI).
- Trike is an open-source threat modeling approach and tool that focuses on using threat models as a risk management tool.
- Construct a platform for Risk Analysis of Security Critical Systems (CORAS), also open source, is a European project that relies heavily on Unified Modeling Language (UML) as the front end for visualizing the threats.
- Visual, Agile, and Simple Threat Modeling (VAST) is a proprietary approach that leverages Agile concepts.

Implementing a structured threat modeling program allows an organization to consistently identify and characterize the threats it faces and then apply appropriate control to the risks associated with those threats.

APPLY SUPPLY CHAIN RISK MANAGEMENT CONCEPTS

The interconnected nature of today's information systems places a high degree of reliance on the confidentiality, integrity, and availability of systems from multiple vendors spread across the globe. This ecosystem has been shown to be vulnerable to both accidental and intentional disruption and compromise. Securing your organization's assets requires that you evaluate the security risk of your entire supply chain and that you apply appropriate controls to manage that risk.

Risks Associated with Hardware, Software, and Services

Any time an organization considers using third-party hardware, software, or services, the organization must determine how the new hardware, software, or services may fit into the organization's existing environment, and evaluate how the additions may impact the organization's overall security posture. For example, if your organization considers using a public cloud provider, there may be compliance risks if the CSP stores data outside of your country, or other security risks if the CSP does not meet data security requirements that you are legally or contractually required to meet.

✔ Malicious Code in the Supply Chain

The widespread use of proprietary commercial off-the-shelf (COTS) software requires customers to trust the security practices of the vendors. However, many instances have been documented where that trust has been abused, and the COTS vendors become a vehicle to introduce vulnerabilities or compromise the CIA aspects of the customers' data.

This method has become increasingly popular for malware authors precisely because the updates are from a trusted source. In 2017, the developer of the antivirus product CCleaner distributed a routine update to its users that contained a remote-access Trojan. As the malicious software had been inserted into the code before it was signed, the entire update package was seen by most users as a legitimate update. More than 2 billion downloads of the compromised software were reported.

✔ SolarWinds and the SUNBURST Attack

One of the largest supply chain attacks in history became public in 2020 when FireEye disclosed a global attack, now known as the SUNBURST attack. SUNBURST is a vulnerability within the SolarWinds Orion Platform, which, if present and activated, allows an attacker to compromise the server on which the Orion product is running.

This widespread attack is particularly concerning because it impacted a SolarWinds product that is used for IT monitoring and management. What should be used to keep an eye on IT infrastructures ironically became the instrument of harm to those infrastructures.

With the SolarWinds Orion product being used by companies around the globe, large and small, this is a devasting example of how important supply chain management is. The victims of the SUNBURST attack include sophisticated tech companies, like Microsoft and Intel, numerous U.S. government agencies, and even the top-tier cybersecurity firm, FireEye (who initially disclosed the breach). In all, SolarWinds estimates that approximately 18,000 firms were affected around the world.

Third-Party Assessment and Monitoring

To minimize supply chain risk, appropriate controls must be applied to verify the security practices of all involved parties. In most cases, controls have been identified

that would address security risks; the toughest challenge is to ensure that third parties actually do what they should to protect your organization's information from those risks.

Any organization that does business with contractors, vendors, or any other third parties should have a third-party risk management policy that establishes a third-party risk management program responsible for assessing, monitoring, and controlling risks associated with outsourcing to third parties. Governance and oversight activities should include onsite security surveys, formal security audits of third-party systems, and penetration testing, where feasible. Any new third party should be assessed against your organization's security requirements, and gaps should be documented and closely monitored. Further, vendors and other third parties should be regularly reassessed and continuously monitored to ensure that they continue to adequately protect your organization's information. We cover audits, audit standards, and other related concepts in detail in Chapter 6.

Minimum Security Requirements

Similar to baselines and standards (discussed earlier in this chapter), your organization should establish minimum security requirements (MSRs) that define the least acceptable security standards that vendors and other parties in your supply chain must satisfy. Of course, you should strive to ensure that your third parties have the strongest possible security postures, but MSRs, as the name suggests, describe the lowest level of security that your organization is willing to accept from a third party. To avoid issues, your MSRs should take into consideration any legal, contractual, or regulatory requirements that you are required to satisfy; you should not establish an MSR that is below any external security compliance requirement. You must also be prepared to audit and assess third parties' compliance with any MSRs that you have established and communicated.

Service-Level Requirements

A *service-level agreement* (*SLA*) is a contractual agreement between a service provider and its customers that establishes the minimum performance standards that the provider is obligated to meet. When dealing with vendors and other third parties, SLAs serve as documented and agreed-upon performance requirements that a customer can use to hold the third party accountable. For example, you may have an SLA with a public cloud provider that commits to a certain level of system uptime and availability. In the event of a sustained outage of the cloud service, you may be entitled to financial compensation or the right to terminate services with no penalty.

Frameworks

Several frameworks explicitly address supply chain risks. This is an evolving area of risk management, but the complexities of managing the information systems supply chain have been evident for many years.

NIST IR 7622

The U.S. government began directly addressing cyber supply chain risk as a separate issue with the publication of NIST IR 7622, "Notional Supply Chain Risk Management Practices for Federal Information Systems." This work recognizes that the actions required of the entities in the supply chain will change depending on their role, as will the level and type of control to be applied. The document identifies 10 practices that should be taken into account in addressing supply chain risk:

- Uniquely identify supply chain elements, processes, and actors.
- Limit access and exposure within the supply chain.
- Establish and maintain the provenance of elements, processes, tools, and data.
- Share information within strict limits.
- Perform supply chain risk management awareness and training.
- Use defensive design for systems, elements, and processes.
- Perform continuous integrator review.
- Strengthen delivery mechanisms.
- Assure sustainment activities and processes.
- Manage disposal and final disposition activities throughout the system or element lifecycle.

The U.S. government has a number of other supply chain risk management initiatives, including the Committee on National Security Systems Directive 505, "Supply Chain Risk Management," which specifically addresses security requirements for strategic national systems and the Comprehensive National Cybersecurity Initiative Number 11, which provides a set of tools to agencies to manage their cybersecurity supply chain through a risk-driven approach.

ISO 28000

ISO 28000:2007, "Specification for security management systems for the supply chain," provides a broad framework for managing supply chain risk. While not specific to cybersecurity, ISO 28000 is useful for organizations that leverage other ISO

specifications (such as ISO 9001 and ISO 27001) to align supply chain risk with the organizations' audit processes or that seek to use a standardized, risk-based approach to evaluating supply chain risk.

ISO 28000:2007 relies heavily on the continuous process improvement model of plan, do, check, act (PDCA) to improve the security management system and to assure organizational conformance to the security practice. This approach facilitates the integration of supply chain risk with broader organizational risk management activities.

U.K. National Cyber Security Centre

The U.K. National Cyber Security Centre (NCSC) proposed guidance that attempts to provide organizations with improved awareness of supply chain risks, while also establishing 12 principles intended to help organizations establish and maintain effective control of their supply chain. The 12 supply chain principles are divided into these separate stages:

1. **Understand the risks:** The principles in this stage involve identifying your vendors in your supply chain and establishing what needs to be protected in that supply chain (and why).
2. **Establish control:** This stage involves establishing minimum security requirements (see the earlier section "Minimum Security Requirements") and communicating your security expectations to your suppliers.
3. **Check your arrangements:** This stage involves establishing assurance activities and building those into your supply chain processes. This includes establishing audit rights, key performance indicators, and other testing/validation activities.
4. **Continuous improvement:** This stage involves continually building trust with your suppliers and constantly encouraging security improvements for your supply chain.

ESTABLISH AND MAINTAIN A SECURITY AWARENESS, EDUCATION, AND TRAINING PROGRAM

No matter how many security tools you have in your arsenal, your organization's security is only as strong as its weakest link — and that tends to be your personnel. Information security is one of the few fields that is governed by relatively small teams but is the responsibility of every person within an organization. As such, all personnel within an organization need to be trained and made aware of security threats and attacker techniques so that they know what to look for and how to avoid common pitfalls that can compromise your organization's information security.

Methods and Techniques to Present Awareness and Training

A *security awareness program* is a formal program that includes processes to train users of the potential threats to an organization's information and systems, as well as educates those users on how to handle such threats. A standard security awareness program should include, at a minimum, new user orientation, lectures or computer-based trainings (CBTs), and printed materials like posters and handouts that share security tips. In addition, organizations can use phishing and other social engineering exercises, security champions, and gamification to help raise awareness of important security topics; each of these is discussed in the following sections.

Social Engineering

Social engineering is the practice of human manipulation that involves an attacker pretending to be someone else in an effort to retrieve sensitive data. *Phishing* is the most common form of social engineering, and it relates to social engineering activities that are conducted over email. Phishing is routinely at the top of the most common security concerns because it can evade many of your most sophisticated security tools and compromise an organization's weakest link — its people.

Simulated phishing campaigns are a popular component of security awareness programs. You should first start by educating your employees on why phishing is harmful and how to spot it. You should conduct randomized simulated phishing exercises to help reinforce the employee training and to help you understand where your risks are (i.e., which types of phishing are most successful on your employees and which employees need further training). Employees who click on a simulated phishing link should be notified and subject to further training that reminds them of how to identify and report signs of phishing.

Security Champions

A *security champion* is a liaison between an organization's security team and the rest of the company; they are tasked with raising security awareness within the organization. In this role, a security champion is an advocate of security best practices for employees who don't work on security as their primary job. The role of security champion was initially created to raise awareness of application security on software development teams, but nowadays, organizations may frequently choose to assign a security champion to any (or all) nonsecurity teams.

Gamification

Gamification is the use of game techniques in nongame applications to engage and educate an audience. In security awareness, gamification can provide a fun and engaging way

to educate employees and promote strong security practices. Games like this allow companies to educate their employees on critical information security concepts in an interesting and engaging manner.

Periodic Content Reviews

Information security is a constantly evolving field, with security threats and vulnerabilities that are forever changing. As such, it's important that you regularly review the content within your security awareness, education, and training program to certify that it remains relevant. Content should be reviewed and updated annually, at a minimum, to ensure that there is no reference to obsolete or irrelevant technologies or terminology, and these reviews should validate that all security awareness and training materials reflect current security trends, concepts, and concerns that are relevant to your organization and industry. Ideally, security awareness content should be considered "live" material that evolves even more frequently than these periodic reviews. As a CISSP, you should ensure that such security training content includes all the relevant and current information that your organization's employees should know.

Program Effectiveness Evaluation

Conducting security awareness, education, and training activities is not enough; it's equally important to evaluate and measure the effectiveness of your security education activities. Although the effectiveness of your security awareness program may be gleaned through the evaluation of your organization's overall information security posture, a formal evaluation should be conducted to target deficiencies within the awareness program itself.

There are several methods by which you can evaluate the effectiveness of your security awareness program. Some examples include the following:

- **Training metrics:** Simple metrics like training completion rates are a great place to start when evaluating the effectiveness of your security awareness program. These types of metrics can tell you whether your training resources are reaching a sufficient percentage of your employees and may alert you if alternate delivery methods are necessary.
- **Quizzes:** This is one of the most effective methods of measuring program effectiveness through knowledge retention. Quizzes are most reliable when measuring the effectiveness of security policies and related information. Analysis of quiz results should be conducted to identify trends that reveal necessary modifications to your training materials; if a substantial number of your employees get the same question wrong, it likely means you need to provide further (or clearer) information about that topic.

- **Security awareness days or weeks:** By sponsoring security awareness days or weeks, you not only have an opportunity to provide security education, but you can also use this as an opportunity to solicit feedback from your employees on the program itself. You can provide attendees with anonymous questionnaires that allow them to express their opinion about the current program and propose new ideas on content delivery.
- **Inherent evaluation:** As previously stated, you can also measure the effectiveness of your awareness program by evaluating your organization's overall security posture. Certain metrics, such as the number of phishing emails and other security issues reported to IT, can provide a great deal of insight into the effectiveness of your program. As your company's employees are increasingly educated on security risks, you should start to see the number of self-reported security issues rise. It's better to see a rise in reported suspected issues than a rise in successful compromises.

SUMMARY

The breadth of information security demands that security professionals possess a wide range of knowledge and skills. You must fully grasp concepts such as confidentiality, integrity, and availability, and understand how to develop, document, and implement security policies, standards, procedures, and guidelines that enforce these concepts. Good security practices must be aligned with an organization's business objectives, strategy, and goals. As a security professional, it's important that you fully understand these business concepts and grasp how you can apply security governance principles to help your organization achieve its mission.

Risk management is at the heart of information security, and every security program should strive to be based on risk management concepts. Identifying threats and vulnerabilities and evaluating security risks is the key to identifying the right security controls to implement in your environment. Controls should be continuously monitored for their effectiveness at reducing risk, and your organization should maintain a program to regularly measure and report on the company's risk posture. There are several industry-standard risk frameworks available to guide your development and management of a risk-based security program.

Legal, regulatory, and compliance requirements play a big role in security. An important component of the CISSP CBK revolves around understanding such laws and other requirements that impact your organization, based on jurisdiction, industry, or other factors.

Tools and technologies often get the majority of the attention, but it's essential to remember that people are the weakest link when it comes to maintaining information security. Candidate screening, background investigations, and other personnel security policies play a critical part in ensuring that your organization's data stays in the right hands. Further, establishing and maintaining a robust security awareness program is an essential way to educate your employees and measure the risk they present to your organization's assets.