CHAPTER 1

SSCP®

Security Operations and Administration

THIS IS WHERE THE planning hits reality; it's in the day to day of information security operations that you see every decision made during the threat assessments and the risk mitigation plans being live-fire tested by your co-workers, customers, legitimate visitors, and threat actors alike. Whether you're an on-shift watch-stander in a security operations center (SOC) or network operations center (NOC) or you work a pattern of normal business hours and days, you'll be exposed to the details of information security in action.

Security operations and administration entail a wide breadth of tasks and functions, and the security professional is expected to have a working familiarity with each of them. This can include maintaining a secure environment for business functions and the physical security of a campus and, specifically, the data center. Throughout your career, you will likely have to oversee and participate in incident response activities, which will include conducting investigations, handling material that may be used as evidence in criminal prosecution and/or civil suits, and performing forensic analysis. The Systems Security Certified Practitioner (SSCP) should also be familiar with common

tools for mitigating, detecting, and responding to threats and attacks; this includes knowledge of the importance and use of event logging as a means to enhance security efforts. Another facet the security practitioner may have to manage could be how the organization deals with emergencies, including disaster recovery.

There is a common thread running through all aspects of this topic: supporting business functions by incorporating security policy and practices with normal daily activities. This involves maintaining an accurate and detailed asset inventory, tracking the security posture and readiness of information technology (IT) assets through the use of configuration/change management, and ensuring personnel are trained and given adequate support for their own safety and security.

This chapter will address all these aspects of security operations. The practitioner is advised, however, to not see this as a thorough treatment of all these concepts, each of which could be (and has been) the subject of an entire book (or books) by themselves; for each topic that is unfamiliar, you should look at the following content as an introduction only and pursue a more detailed review of related subject matter.

NOTE The countries and regions that an organization operates in may have varying, distinct, and at times conflicting legal systems. Beyond considerations of written laws and regulations, the active functioning of court systems and regulatory bodies often has intricate, myriad applications in the real world that extend far beyond how things are codified in written laws. These factors become even more varied and complex when an organization functions in multiple countries and needs to deal with actual scenarios that directly involve international law and the laws of each respective nation. With that in mind, it is always imperative to get the input of a professional legal team to fully understand the legal scope and ramifications of security operations (and basically all operations and responsibilities beyond security as well).

COMPLY WITH CODES OF ETHICS

Your day-to-day journey along the roadmap of security operations and administration must keep one central ideal clearly in focus. Every day that you serve as an information security professional, you make or influence decisions. Every one of those decision moments is an opportunity or a vulnerability; it is a moment in which you can choose to

do the technically and ethically correct thing or the expedient thing. Each of those decision moments is a test for you.

Those decisions must be ethically sound; yes, they must be technically correct, cost-effective, and compliant with legal and regulatory requirements, but at their heart they must be *ethical*. Failure to do so puts your professional and personal integrity at risk, as much as it puts your employer's or your clients' reputation and integrity at risk.

Being a security professional requires you to work, act, and think in ways that comply with and support the codes of ethics that are fundamental parts of your workplace, your profession, and your society and culture at large. Those codes of ethics should harmonize with if not *be* the fundamental ethical values and principles you live your life by—if they do not, that internal conflict in values may make it difficult if not impossible to achieve a sense of personal *and* professional integrity! Professional and personal integrity should be wonderfully, mutually self-reinforcing.

Let's first focus on what ethical decision-making means. This provides a context for how you, as an SSCP, comply with and support the (ISC)^2 Code of Ethics in your daily work and life. We'll see that this is critical to being able to live up to and fulfill the "three dues" of your responsibilities: due care, due diligence, and due process.

Understand, Adhere to, and Promote Professional Ethics

Let's start with what it means to be a professional: It means that society has placed great trust and confidence in you, because you have been willing to take on the responsibility to get things done right. Society trusts in you to know your practice, know its practical limits, and work to make sure that the services you perform meet or exceed the best practices of the profession. This is a legal and an ethical responsibility.

Everything you do requires you to understand the needs of your employers or clients. You listen, observe, gather data, and ask questions; you think about what you've learned, and you come to conclusions. You make recommendations, offer advice, or take action within the scope of your job and responsibilities. Sometimes you take action outside of that scope, going above and beyond the call of those duties. You do this because you are a professional. You would not even think of making those conclusions or taking those actions if they violently conflicted with what known technical standards or recognized best technical practice said was required. You would not knowingly recommend or act to violate the law. Your professional ethics are no different. They are a set of standards that are both constraints and freedoms that you use to inform, shape, and then test your conclusions and decisions with before you act.

As a professional—in any profession—you learned what that profession requires of you through education, training, and on-the-job experience. You learned from teachers, mentors, trainers, and the people working alongside of you. They shared their hard-earned insight and knowledge with you, as their part of promoting the profession you had

in common. In doing so they strengthened the practice of the ethics of the profession, as well as the practice of its technical disciplines.

(ISC)² Code of Ethics

(ISC)² provides a Code of Ethics, and to be an SSCP, you agree to abide by it. It is short and simple. It starts with a preamble, which is quoted here in its entirety:

> *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 is a condition of certification.*

Let's operationalize that preamble—take it apart, step-by-step, and see what it really asks of us.

- **Safety and welfare of society:** Allowing information systems to come to harm because of the failure of their security systems or controls can lead to damage to property or injury or death of people who were depending upon those systems operating correctly.
- **The common good:** All of us benefit when our critical infrastructures, providing common services that we all depend upon, work correctly and reliably.
- **Duty to our principals:** Our duties to those we regard as leaders, rulers, or our supervisors in any capacity.
- **Our duty to each other:** To our fellow SSCPs, others in our profession, and to others in our neighborhood and society at large.
- **Adhere and be seen to adhere to:** Behave correctly and set the example for others to follow. Be visible in performing your job ethically (in adherence with this code) so that others can have confidence in us as a profession and learn from our example.

The code is equally short, containing just four canons or principles to abide by.

> *Protect society, the common good, necessary public trust and confidence, and the infrastructure.*
>
> *Act honorably, honestly, justly, responsibly, and legally.*
>
> *Provide diligent and competent service to principals.*
>
> *Advance and protect the profession.*

The canons do more than just restate the preamble's two points. They show you *how* to adhere to the preamble. You must take action to protect what you value; that action should be done with honor, honesty, and with justice as your guide. Due care and due diligence are what you owe to those you work for (including the customers of the businesses that employ us!).

The final canon talks to your continued responsibility to grow as a professional. You are on a never-ending journey of learning and discovery; each day brings an opportunity to make the profession of information security stronger and more effective. You as an SSCP are a member of a worldwide *community of practice*—the informal grouping of people concerned with the safety, security, and reliability of information systems and the information infrastructures of the modern world.

In ancient history, there were only three professions—those of medicine, the military, and the clergy. Each had in its own way the power of life and death of individuals or societies in its hands. Each as a result had a significant burden to be the best at fulfilling the duties of that profession. Individuals felt the calling to fulfill a sense of duty and service, to something larger than themselves, and responded to that calling by becoming a member of a profession.

This, too, is part of being an SSCP. Visit `https://www.isc2.org` for more information.

Organizational Code of Ethics

Most businesses and nonprofit or other types of organizations have a code of ethics that they use to shape their policies and guide them in making decisions, setting goals, and taking actions. They also use these codes of ethics to guide the efforts of their employees, team members, and associates; in many cases, these codes can be the basis of decisions to admonish, discipline, or terminate their relationship with an employee. In most cases, organizational codes of ethics are also extended to the partners, customers, or clients that the organization chooses to do business with. Sometimes expressed as *values* or *statements of principles*, these codes of ethics may be in written form, established as policy directives upon all who work there; sometimes, they are implicitly or tacitly understood as part of the organizational culture or shaped and driven by key personalities in the organization. But just because they aren't written down doesn't mean that an ethical code or framework for that organization doesn't exist.

Fundamentally, these codes of ethics *have the capacity to* balance the conflicting needs of law and regulation with the bottom-line pressure to survive and flourish as an organization. This is the real purpose of an organizational ethical code. Unfortunately, many organizations let the balance go too far toward the bottom-line set of values and take shortcuts; they compromise their ethics, often end up compromising their legal or regulatory responsibilities, and end up applying their codes of ethics loosely if at all. As a case in point, consider that risk management must include the dilemma that sometimes there are more laws and regulations than any business can possibly afford to comply with *and* they all conflict with each other in some way, shape, or form. What's a chief executive or a board of directors to do in such a circumstance?

It's actually quite easy to incorporate professional and personal ethics, along with the organization's own code of ethics, into every decision process you use. Strengths,

weaknesses, opportunities, and threats (SWOT) analyses, for example, focus your attention on the strengths, weaknesses, opportunities, and threats that a situation or a problem presents; being true to one's ethics should be a *strength* in such a context, and if it starts to be seen as a weakness or a threat, that's a danger signal you must address or take to management and leadership. Cost/benefits analyses or decision trees present the same opportunity to include what sometimes is called the *New York Times* or the *Guardian* test: How would each possible decision look if it appeared as a headline on such newspapers of record? Closer to home, think about the responses you might get if you asked your parents, family, or closest friends for advice about such thorny problems—or their reactions if *they* heard about it via their social media channels. Make these thoughts a habit; that's part of the *practice* aspect of being a professional.

As the on-scene information security professional, you'll be the one who most likely has the first clear opportunity to look at an IT security posture, policy, control, or action, and challenge any aspects of it that you think might conflict with the organization's code of ethics, the $(ISC)^2$ Code of Ethics, or your own personal and professional ethics.

UNDERSTAND SECURITY CONCEPTS

What does it mean to "keep information secure?" What is a good or adequate "security posture?" Let's take questions like these and operationalize them by looking for characteristics or attributes that measure, assess, or reveal the overall security state or condition of our information.

- **Confidentiality:** Limits are placed on who is allowed to view the information, including copying it to another form.
- **Integrity:** The information stays complete and correct when retrieved, displayed, or acted upon.
- **Availability:** The information is presented to the user in a timely manner when required and in a form and format that meets the user's needs.
- **Authenticity:** Only previously approved, known, and trusted users or processes have been able to create, modify, move, or copy the information.
- **Utility:** The content of the information, its form and content, and its presentation or delivery to the user meet the user's needs.
- **Possession or control:** The information is legally owned or held by a known, authorized user, such that the user has authority to exert control over its use, access, modification, or movement.

- **Safety:** The system and its information, by design, do not cause unauthorized harm or damage to others, their property, or their lives.
- **Privacy:** Information that attests to or relates to the identity of a person, or links specific activities to that identity, must be protected from being accessed, viewed, copied, modified, or otherwise used by unauthorized persons or systems.
- **Nonrepudiation:** Users who created, used, viewed, or accessed the information, or shared it with others, cannot later deny that they did so.
- **Transparency:** The information can be reviewed, audited, and made visible or shared with competent authorities for regulatory, legal, or other processes that serve the public good.

Note that these are characteristics of the information itself. Keeping information authentic, for example, levies requirements on all of the business processes and systems that could be used in creating or changing that information or changing anything about the information.

All of these attributes boil down to one thing: *decision assurance*. How much can we trust that the decisions we're about to make are based on reliable, trustworthy information? How confident can we be that the competitive advantage of our trade secrets or the decisions we made in private are still unknown to our competitors or our adversaries? How much can we count on that decision being the right decision, in the legal, moral, or ethical sense of its being correct and in conformance with accepted standards?

Another way to look at attributes like these is to ask about the *quality* of the information. Bad data—data that is incomplete, incorrect, not available, or otherwise untrustworthy—causes monumental losses to businesses around the world; an IBM study reported that in 2017 those losses exceeded $3.1 trillion, which may be more than the total losses to business and society due to information security failures. Paying better attention to a number of those attributes would dramatically improve the reliability and integrity of information used by any organization; as a result, a growing number of information security practitioners are focusing on data quality as something they can contribute to.

Conceptual Models for Information Security

There are any number of frameworks, often represented by their acronyms, which are used throughout the world to talk about information security. All are useful, but some are more useful than others.

- *The CIA triad* (sometimes written as CIA) combines confidentiality, integrity, and availability and dates from work being done in the 1960s to develop theoretical models for information systems security and then implement those technologies into operating systems, applications programs, and communications and network systems.

- *CIANA* combines confidentiality, integrity, availability, nonrepudiation, and authentication. The greater emphasis on nonrepudiation and authentication provides a much stronger foundation for both criminal and civil law to be able to ascertain what actions were taken, by whom, and when, in the context of an incident, dispute, or conflicting claims of ownership or authorship.
- CIANA+PS expands CIANA to include privacy and safety. Cyberattacks in the Ukraine since 2014 and throughout the world from 2017 to present highlight the need for far more robust operational technology (OT) safety and resiliency. At the same time, regulators and legislators continue to raise the standards for protecting privacy-related data about individuals, with over 140 countries having privacy data protection laws in effect.
- *The Parkerian hexad* includes confidentiality, integrity, availability, authenticity, utility, and possession or control.

These frameworks, and many more, have their advocates, their user base, and their value. That said, in the interest of consistency, we'll focus throughout this book on CIANA+PS, as its emphasis on both nonrepudiation and authentication have perhaps the strongest and most obvious connections to the vitally important needs of e-commerce and our e-society to be able to conduct personal activities, private business, and governance activities in ways that are safe, respectful of individual rights, responsible, trustworthy, reliable, and transparent.

It's important to keep in mind that these attributes of systems performance or effectiveness build upon each other to produce the overall degree of trust and confidence we can rightly place on those systems and the information they produce for us. We *rely* on high-reliability systems because their information is correct and complete (high integrity), it's where we need it when we need it (availability), and we know it's been kept safe from unauthorized disclosure (it has authentic confidentiality), while at the same time we have confidence that the only processes or people who've created or modified it are trusted ones. Our whole sense of "can we trust the system and what it's telling us" is a greater conclusion than just the sum of the individual CIANA+PS, Parkerian, or triad attributes.

Let's look further at some of these attributes of information security.

Confidentiality

Often thought of as "keeping secrets," confidentiality is actually about sharing secrets. Confidentiality is both a legal and ethical concept about *privileged communications* or *privileged information*. Privileged information is information you have, own, or create, and that you share with someone else with the agreement that they cannot share that knowledge with anyone else without your consent or without due process in law. You place your trust and confidence in that other person's adherence to that agreement. Relationships between professionals and their clients, such as the doctor-patient or attorney-client ones,

are prime examples of this privilege in action. In rare exceptions, courts cannot compel parties in a privileged relationship to violate that privilege and disclose what was shared in confidence.

Confidentiality refers to how much we can trust that the information we're about to use to make a decision with has not been seen by unauthorized people. The term *unauthorized people* generally refers to any person or any group of people who could learn something from our confidential information and then use that new knowledge in ways that would thwart our plans to attain our objectives or cause us other harm.

Confidentiality needs dictate who can read specific information or files or who can download or copy them; this is significantly different from who can modify, create, or delete those files.

One way to think about this is that integrity violations change what *we* think *we* know; confidentiality violations tell others what *we* think is *our* private knowledge.

Business has many categories of information and ideas that it needs to treat as confidential, such as the following:

- Proprietary, or company-owned information, whether or not protected by patent, copyright, or trade secret laws
- Proprietary or confidential information belonging to others but shared with the company under the terms of a nondisclosure agreement (NDA)
- Company private data, which can include business plans, budgets, risk assessments, and even organizational directories and alignments of people to responsibilities
- Data required by law or regulation to be kept private or confidential
- Privacy-related information pertaining to individual employees, customers, prospective customers or employees, or members of the public who contact the firm for any reason
- Customer transaction and business history data, including the company's credit ratings and terms for a given customer
- Customer complaints, service requests, or suggestions for product or service improvements

In many respects, such *business confidential* information either represents the results of investments the organization has already made or provides insight that informs decisions they're about to make; either way, all of this and more represent *competitive advantage* to the company. Letting this information be disclosed to unauthorized persons, *inside or outside* of the right circles within the company, threatens to reduce the value of those investments and the future return on those investments. It could, in the extreme, put the company out of business!

Let's look a bit closer at how to defend such information.

Intellectual Property

Our intellectual property are the ideas that we create and express in tangible, explicit form; in creating them, we create an ownership interest. Legal and ethical frameworks have long recognized that such creativity benefits a society and that such creativity needs to be encouraged and incentivized. Incentives can include financial reward, recognition and acclaim, or a legally protected ownership interest in the expression of that idea and its subsequent use by others. This vested interest was first recognized by Roman law nearly 2,000 years ago. Recognition is a powerful incentive to the creative mind, as the example of the Pythagorean theorem illustrates. It was created long before the concept of patents, rights, or royalties for intellectual property were established, and its creator has certainly been dead for a long time, and yet no ethical person would think to attempt to claim it as their own idea. Having the author's name on the cover of a book or at the masthead of a blog post or article also helps to recognize creativity.

Financial reward for ideas can take many forms, and ideally, such ideas should pay their own way by generating income for the creator of the idea, recouping the expenses they incurred to create it, or both. Sponsorship, grants, or the salary associated with a job can provide this; creators can also be awarded prizes, such as the Nobel Prize, as both recognition and financial rewards.

The best incentive for creativity, especially for corporate-sponsored creativity, is in how that ownership interest in the new idea can be turned into profitable new lines of business or into new products and services.

The vast majority of intellectual property is created in part by the significant investment of private businesses and universities in both basic research and product-focused developmental research. Legal protections for the intellectual property (or IP) thus created serve two main purposes. The first is to provide a limited period of time in which the owner of that IP has a monopoly for the commercial use of that idea and thus a sole claim on any income earned by selling products or providing services based on that idea. These monopolies were created by an edict of the government or the ruling monarchy, with the first being issued by the Doge of Venice in the year 1421. Since then, nation after nation has created patent law as the body of legal structure and regulation for establishing, controlling, and limiting the use of patents. The monopoly granted by a patent is limited in time and may even (based on applicable patent law) be limited in geographic scope or the technical or market reach of the idea. An idea protected by a patent issued in Colombia, for example, may not enjoy the same protection in Asian markets as an idea protected by U.S., U.K., European Union, or Canadian patent law. The second purpose is to publish the idea itself to the marketplace so as to stimulate rapid adoption of the idea, leading to widespread adoption, use, and influence upon the marketplace and upon society. Patents may be *monetized* by selling the rights to the patent or by licensing the use of the patent to another person or business; income from such licensing or sale has

long been called the *royalties* from the patent (in recognition that it used to take an act of a king or a queen to make a patent enforceable).

Besides patents and patent law, there exist bodies of law regarding copyrights, trademarks, and trade secrets. Each of these treats the fruits of one's intellectually creative labors differently, and like patent law, these legal and ethical constructs are constantly under review by the courts and the cultures they apply to. Patents protect an idea, a process, or a procedure for accomplishing a practical task. Copyrights protect an artistic expression of an idea, such as a poem, a painting, a photograph, or a written work (such as this book). Trademarks identify an organization or company and its products or services, typically with a symbol, an acronym, a logo, or even a caricature or character (not necessarily of a person). Trade secrets are the unpublished ideas, typically about step-by-step details of a process, or the recipe for a sauce, paint, pigment, alloy, or coating, that a company or individual has developed. Each of these represent a competitive advantage worthy of protection. Note the contrast in these forms, as shown in Table 1.1.

TABLE 1.1 Forms of Intellectual Property Protection

LEGAL CONCEPT	PUBLIC DISCLOSURE	MONETIZE BY	COMPROMISE BY
Patent	Mandatory, detailed	License to use	Failure to develop or monetize; failure to defend against infringement
Copyright	Published works	Sell copies	Failure to defend
Trademark	Logos, signs, product stampings	Creates brand awareness in marketplace	Failure to defend
Trade secret	Must be undisclosed	Sell products and services based on its use; can be licensed	Failure to keep secret or defend

The most important aspect of that table for you, as the on-scene information security professional, is the fourth column. Failure to defend and failure to keep secret both require that the owners and licensed or authorized users of a piece of IP must take all reasonable, prudent efforts to keep the ideas and their expression in tangible form safe from infringement. This protection must be firmly in place throughout the entire lifecycle of the idea—from its first rough draft of a sketch on the back of a cocktail napkin through drawings, blueprints, mathematical models, and computer-aided design and manufacturing (CADAM) data sets. All expressions of that idea in written, oral, digital, or physical form must be protected from inadvertent disclosure or deliberate but unauthorized viewing or copying. Breaking this chain of confidentiality can lead to voiding the claim

to protection by means of patent, copyright, or trade secret law. In its broadest terms, this means that the organization's information systems must ensure the confidentiality of this information.

Protect IP by Labeling It

Protection of intellectual property must consider three possible exposures to loss: exfiltration, inadvertent disclosure, and failure to aggressively assert one's claims to protection and compensation. Each of these is a failure by the organization's management and leadership to exercise due care and due diligence.

- *Exfiltration* generally occurs in part because decisions have been made to ignore risks, disregard alarm indications, and knowingly operate information systems in insecure ways. (There are cases of data breaches that happen to highly secure systems, hardened to the best possible standards, but these are few and far between.)
- Inadvertent exposure can happen due to carelessness, due to accident, or through faulty design of business processes or information security measures.
- An expression of an idea must, in almost all cases, be labeled or declared as a protected idea; this is how its owner asserts rights against possible infringement. This first assertion of a claim of ownership provides the basis for seeking legal means to stop the infringement, seek damages for lost business, or enter into licensing arrangements with the infringers.

Each of these possible exposures to loss starts with taking proper care of the data in the first place. This requires properly classifying it (in terms of the restrictions on handling, use, storage, or dissemination required), marking or labeling it (in human-readable and machine-readable ways), and then instituting procedures that enforce those restrictions.

Software, Digital Expression, and Copyright

Most software is protected by copyright, although a number of important software products and systems are protected by patents. Regardless of the protection used, it is implemented via a license. Most commercially available software is not actually sold; customers purchase a license to use it, and that license strictly limits that use. This license for use concept also applies to other copyrighted works, such as books, music, movies, or other multimedia products. As a customer, you purchase a license for its use (and pay a few pennies for the DVD, Blu-Ray, or other media it is inscribed upon for your use). In most cases, that license prohibits you from making copies of that work and from giving copies to others to use. Quite often, they are packaged with a copy-protection feature or with features that engage with digital rights management (DRM) software that is increasingly part of modern operating systems, media player software applications, and home

and professional entertainment systems. It is interesting to note that, on one hand, digital copyright law authorizes end-user licensees to make suitable copies of a work in the normal course of consuming it—you can make backups, for example, or shift your use of the work to another device or another moment in time. On the other hand, the same laws prohibit you from using any reverse engineering, tools, processes, or programs to defeat, break, side-step, or circumvent such copy protection mechanisms. Some of these laws go further and specify that attempts to defeat any encryption used in these copy protection and rights management processes is a separate crime itself.

These laws are part of why businesses and organizations need to have acceptable use policies in force that control the use of company-provided IT systems to install, use, consume, or modify materials protected by DRM or copy-protect technologies. The employer, after all, can be held liable for damages if they do not exert effective due diligence in this regard and allow employees to misuse their systems in this way.

Copyleft?

By contrast, consider the Creative Commons license, sometimes referred to as a copyleft. The creator of a piece of intellectual property can choose to make it available under a Creative Commons license, which allows anyone to freely use the ideas provided by that license so long as the user attributes the creation of the ideas to the licensor (the owner and issuer of the license). Businesses can choose to share their intellectual property with other businesses, organizations, or individuals by means of licensing arrangements. Copyleft provides the opportunity to widely distribute an idea or a practice and, with some forethought, leads to creating a significant market share for products and services. Pixar Studios, for example, has made RenderMan, its incredibly powerful, industry-leading animation rendering software, available free of charge under a free-to-use license that is a variation of a creative commons license. In March 2019, the National Security Agency made its malware reverse engineering software, called Ghidra, publicly available (and has since issued bug fix releases to it). Both approaches reflect a savvy strategy to influence the ways in which the development of talent, ideas, and other products will happen in their respective marketplaces.

Industrial or Corporate Espionage

Corporations constantly research the capabilities of their competitors to identify new opportunities, technologies, and markets. Market research and all forms of open source intelligence (OSINT) gathering are legal and ethical practices for companies, organizations, and individuals to engage in. Unfortunately, some corporate actors extend their research beyond the usual venue of trade shows and reviewing press releases and seek to conduct surveillance and gather intelligence on their competitors in ways that move along the ethical continuum from appropriate to unethical and, in some cases, into

illegal actions. In many legal systems, such activities are known as *espionage*, rather than research or business intelligence, as a way to clearly focus on their potentially criminal nature. (Most nations consider it an illegal violation of their sovereignty to have another nation conduct espionage operations against it; most nations, of course, conduct espionage upon each other regardless.) To complicate things even further, nearly all nations actively encourage their corporate citizens to gather business intelligence information about the overseas markets they do business in, as well as about their foreign or multinational competitors operating in their home territories. The boundary between corporate espionage and national intelligence services has always been a blurry frontier.

When directed against a competitor or a company trying to enter the marketplace, corporate-level espionage activities that might cross over an ethical or legal boundary can include attempts to do the following:

- Establish business relationships to gain federated access to e-business information such as catalogs, price lists, and specifications
- Gather product service or maintenance manuals and data
- Recruit key personnel from the firm, either as new employees or as consultants
- Engaging in competitive, information-seeking arrangements with key suppliers, service vendors, or customers of the target firm
- Probing and penetration efforts against the target's websites and online presence
- Social engineering efforts to gather intelligence data or provide the reconnaissance footprint for subsequent data gathering
- Unauthorized entry or breaking into the target's property, facilities, or systems
- Visiting company facilities or property, ostensibly for business purposes, but as intelligence-gathering

All of the social engineering techniques used by hackers and the whole arsenal of advanced persistent threat (APT) tools and techniques might be used as part of an industrial espionage campaign. Any or all of these techniques can and often are done by third parties, such as hackers (or even adolescents), often through other intermediaries, as a way of maintaining a degree of plausible deniability.

You will probably never know if that probing and scanning hitting your systems today has anything to do with the social engineering attempts by phone or email of a few weeks ago. You'll probably never know if they're related to an industrial espionage attempt, to a ransom or ransomware attack, or as part of an APT's efforts to subvert some of your systems as launching pads for their attacks on other targets. Protect your systems against each such threat vector as if each system does have the defense of the company's

intellectual property "crown jewels" as part of its mission. That's what *keeping confidences*, what protecting the confidential, proprietary, or business-private information, comes down to, doesn't it?

Integrity

Integrity, in the common sense of the word, means that something is whole, complete, its parts smoothly joined together. People with high personal integrity are ones whose actions and words consistently demonstrate the same set of ethical principles. Having such integrity, you know you can count on them and trust them to act both in ways they have told you they would and in ways consistent with what they've done before.

When talking about information systems, *integrity* refers to both the information in them and the processes (that are integral to that system) that provide the functions we perform on that information. Both of these—the information and the processes—must be complete, correct, function together correctly, and do so in reliable, repeatable, and deterministic ways for the overall system to have integrity.

When we measure or assess information systems integrity, therefore, we can think of it in several ways.

- **Binary:** Either our information system has integrity or it does not. We can rely upon it or we cannot.
- **Threshold-based:** Our information system has at least a minimum level of systems and information integrity to function reliably but possibly in a degraded way, either with higher than desired (but still acceptable) error rates or at reduced transaction throughput or volume levels.

Note that in all but the simplest of business or organizational architectures, you'll find multiple sets of business logic and therefore business processes that interact with each other throughout overlapping cycles of processing. Some of these *lines of business* can function independently of each other, for a while, so long as the information and information systems that serve that line of business directly are working correctly (that is, have high enough levels of integrity).

- Retail online sales systems have customer-facing processes to inform customers about products, services, and special offers. Their shopping cart systems interact with merchandise catalog databases, as well as with order completion, payment processing, and order fulfillment. Customer sales order processing and fulfillment can occur—with high integrity—even though other systems that update the catalogs to reflect new products or services or bring new vendors and new product lines into the online store are not available.

- Computer-aided manufacturing systems have to control the flow of materials, parts, subassemblies, and finished products on the factory floor, interacting with logistics and warehousing functions on both the input and output sides of the assembly line. These systems are typically not tightly coupled with the functions of other business elements, such as finance, sales and marketing, or personnel management, even though at some point the assembly line grinds to a halt if finance hasn't paid the bills to suppliers in a timely way.

REAL WORLD EXAMPLE: Trustworthiness Is Perceptual

You make a decision to trust in what your systems are telling you. You choose to believe what the test results, the outputs of your monitoring systems, and your dashboards and control consoles are presenting to you as "ground truth," the truth you could observe if you were right there on the ground where the event reported by your systems is taking place. Most of the time, you're safe in doing so.

The operators of Iran's nuclear materials processing plant believed what their control systems were reporting to them, all the while the Stuxnet malware had taken control of both the processing equipment and the monitoring and display systems. Those displays lied to their users, while Stuxnet drove the uranium processing systems to self-destruct.

An APT that gets deep into your system can make your systems lie to you as well. Attackers have long used the techniques of perception management to disguise their actions and mislead their targets' defenders.

Your defense: Find a separate and distinct means for verifying what your systems are telling you. Get out-of-band or out-of-channel and gather data in some other way that is as independent as possible from your mainline systems; use this alternative source intelligence as a sanity check.

Integrity applies to three major elements of any information-centric set of processes: to the people who run and use them, to the data that the people need to use, and to the systems or tools that store, retrieve, manipulate, and share that data. Note, too, that many people in the IT and systems world talk about "what we know" in four very different but strongly related ways, sometimes referred to as D-I-K-W.

- Data consists of the individual facts, observations, or elements of a measurement, such as a person's name or their residential address.
- Information results when you process data in various ways; information is data plus conclusions or inferences.

- Knowledge is a set of broader, more general conclusions or principles that you've derived from lots of information.
- Wisdom is (arguably) the insightful application of knowledge; it is the "a-ha!" moment in which you recognize a new and powerful insight that you can apply to solve problems with or take advantage of a new opportunity—or to resist the temptation to try!

Figure 1.1 illustrates this knowledge pyramid.

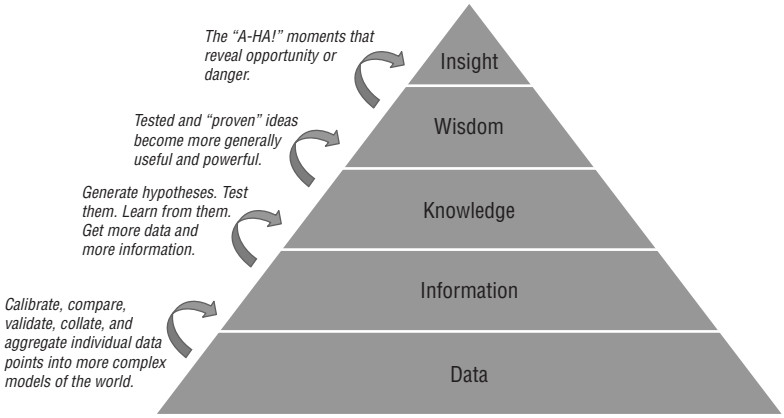

FIGURE 1.1 **The DIKW knowledge pyramid**

Professional opinion in the IT and information systems world is strongly divided about data versus DIKW, with about equal numbers of people holding that they are the same ideas, that they are different, and that the whole debate is unnecessary. As an information security professional, you'll be expected to combine experience, training, and the data you're observing from systems and people in real time to know whether an incident of interest is about to become a security issue, whether your organization uses knowledge management terminology like this or not. This is yet another example of just how many potentially conflicting, fuzzy viewpoints exist in IT and information security.

Availability

Is the data there when we need it in a form we can use?

We make decisions based on information; whether that is new information we have gathered (via our data acquisition systems) or knowledge and information we have in our memory, it's obvious that if the information is not where we need it when we need it, we cannot make as good a decision as we might need.

- The information might be in our files, but if we cannot retrieve it, organize it, and display it in ways that inform the decision, then the information isn't available.
- If the information has been deleted, by accident, sabotage, or systems failure, then it's not available to inform the decision.

Those might seem obvious, and they are. Key to availability requirements is that they specify what information is needed; where it will need to be displayed, presented, or put in front of the decision-makers; and within what span of time the data is both available (displayed to the decision-makers) and meaningful. Yesterday's data may not be what we need to make today's decision.

Note that availability means something different for a system than it does for the information the system produces for us. *Systems availability* is measurable, such as via a percentage of capacity or a throughput rate. *Information availability*, by contrast, tells us one of three things.

- Yes, we have what we need to know to make this decision or take this action.
- No, we do not have what we need to know, so we have to decide blindly.
- We have *some* of what we need to know, and we cannot logically infer that what's missing won't cause our decision to be wrong and lead us to harm.

Accountability

Information and information systems represent significant investments by organizations, and as a result, there's a strong bottom-line financial need to know that such investments are paying off—and that their value is not being diluted due to loss of control of that information (via a data breach or exfiltration) or loss or damage to the data's integrity or utility. Organizations have three functional or operational needs for information regarding accountability. First, they gather information about the *use* of corporate information and IT systems. Then they consolidate, analyze, and audit that usage information. Finally, they use the results of those reviews to inform decision-making. Due diligence needs, for example, are addressed by resource chargeback, which attributes the per-usage costs of information to each internal user organization. Individuals must also be held accountable for their own actions, including their use or misuse of corporate information systems. Surrounding all of this is the need to know whether the organization's information security systems are actually working correctly and that alarms are being properly attended to.

Privacy

Although legal and cultural definitions of privacy abound, we each have an internalized, working idea of what it means to keep something *private*. Fundamentally, this means that when we do something, write something down, or talk with another person, we have a reasonable expectation that what is said and done stays within a space and a place that we can control. *We* get to choose whom we share our thoughts with or whom we invite into our home. And with this working concept of privacy deep in our minds, we establish circles of trust. The innermost circle, those closest to us, we call our *intimates*; these are

the people with whom we mutually share our feelings, ideas, hopes, worries, and dreams. Layer by layer, we add on other members of our extended family, our neighbors, or even people we meet every morning at the bus stop. We know these people to varying degrees, and our trust and confidence in them varies as well. We're willing to let our intimates make value judgments about what we consider to be our private matters or accept criticism from them about such matters; we don't share these with those not in our "inner circle," and we simply not tolerate them (tolerate criticism or judgments) from someone who is not at the same level of trust and regard.

Businesses work the same way. Businesses need to have a reasonable expectation that problems or issues stay within the set of people within the company who need to be aware of them and involved in their resolution. This is in addition to the concept of business confidential or proprietary information—it's the need to take reasonable and prudent measures to keep conversations and tacit knowledge inside the walls of the business and, when applicable, within select circles of people inside the business.

✔ Privacy Is Not Confidentiality

As more and more headline-making data breaches occur, people are demanding greater protection of personally identifiable information (PII) and other information about them as individuals. Increasingly, this is driving governments and information security professionals to see *privacy* as separate and distinct from *confidentiality*. While both involve keeping closely held, limited-distribution information safe from inadvertent disclosure, we're beginning to see that they may each require subtly different approaches to systems design, operation, and management to achieve.

Privacy: In Law, in Practice, in Information Systems

In legal terms, privacy relates to three main principles: restrictions on search and seizure of information and property, self-incrimination, and disclosure of information held by the government to plaintiffs or the public. Many of these legal concepts stem from the idea that government must be restricted from taking arbitrary action against its citizens, or people (human beings or fictitious entities) who are within the jurisdiction of those governments. Laws such as the Fourth and Fifth Amendments to the US Constitution, for example, address the first two, while the Privacy Act of 1974 created restrictions on how government could share with others what it knew about its citizens (and even limited sharing of such information within the government). Medical codes of practice and the laws that reflect them encourage data sharing to help health professionals detect a potential new disease epidemic but also require that personally identifiable information in the clinical data be removed or anonymized to protect individual patients.

The European Union has enacted a series of policies and laws designed to protect individual privacy as businesses and governments exchange data about people, about transactions, and about themselves. The latest of these, the General Data Protection Regulation 2016/679, is a law binding upon all persons, businesses, or organizations doing anything involving the data related to an EU person. GDPR's requirements meant that by May 2018, businesses had to change the ways that they collected, used, stored, and shared information about anyone who contacted them (such as by browsing to their website); they also had to notify such users about the changes and gain their informed consent to such use. Many news and infotainment sites hosted in the United States could not serve EU persons until they implemented changes to become GDPR compliant.

Privacy as a data protection framework, such as GDPR, provides you with specific functional requirements your organization's use of information must comply with; you are a vital part in making that compliance effective and in assuring that such usage can be audited and controlled effectively. If you have doubts as to whether a particular action or an information request is legal or ethical, ask your managers, the organizational legal team, or its ethics advisor (if it has one).

In some jurisdictions and cultures, we speak of an inherent right to privacy; in others, we speak to a requirement that people and organizations protect the information that they gather, use, and maintain when that data is about another person or entity. In both cases, the right or requirement exists to prevent harm to the individual. Loss of control over information about you or about your business can cause you grave if not irreparable harm.

Law at local, national, and international levels continues to evolve. Let's look at a fews.

Universal Declaration of Human Rights

Following World War II, there was a significant renewal and an increased sense of urgency to ensure that governments did not act in an arbitrary manner against citizens. The United Nations drafted the Universal Declaration of Human Rights that set forth these expectations for members. Article 12 states, "No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks."

OECD and Privacy

The Organization for Economic Cooperation and Development (OECD) promotes policies designed to improve the economic and social well-being of people around the world. In 1980, the OECD published "Guidelines on the Protection of Privacy and Transborder Flows of Personal Data" to encourage the adoption of comprehensive privacy protection practices. In 2013, the OECD revised its Privacy Principles to address the wide range of

challenges that came about with the explosive growth of information technology. Among other changes, the new guidelines placed greater emphasis on the role of the data controller to establish appropriate privacy practices for their organizations.

✔ OECD Privacy Principles: Basic Principles of National Application

The OECD Privacy Principles are used throughout many international privacy and data protection laws and are also used in many privacy programs and practices. The eight privacy principles are as follows:

1. **Collection Limitation Principle:** This principle states that data that is collected should be obtained by lawful and fair means, that the data subject should be aware of and consent to the collection of the data where appropriate, and that the quantity and type of data should be limited.
2. **Data Quality Principle:** This principle is aimed at the accuracy and completeness of data, whether it is appropriately maintained and updated, and whether the data retained is relevant to the purposes it is used for.
3. **Purpose Specification Principle:** Purpose specification means that the reasons that personal data is collected should be determined before it is collected, rather than after the fact, and that later data reuse is in line with the reason that the data was originally obtained.
4. **Use Limitation Principle Security:** This principle notes that release or disclosure of personal data should be limited to the purposes it was gathered for unless the data subject agrees to the release or it is required by law.
5. **Security Safeguards Principle:** Reasonable security safeguards aimed at preventing loss, disclosure, exposure, use, or destruction of the covered data are the focus of this principle.
6. **Openness Principle:** The principle of openness is intended to ensure that the practices and policies that cover personal data are accessible and that the existence of personal data, what data is collected and stored, and what it is used for should all be disclosed. Openness also requires that the data controller's identity and operating location or residence is openly disclosed.
7. **Individual Participation Principle:** This includes an individual's right to know if their data has been collected and stored and what that data is within a reasonable time and in a reasonable way. In addition, this principle allows the

CONTINUES

subject to request that the data be corrected, deleted, or otherwise modified as needed. An important element of this principle is the requirement that data controllers must also explain why any denials of these rights are made.

8. **Accountability Principle:** The final principle makes the data controller accountable for meeting these principles.

The OECD Privacy Guidelines can be found at `www.oecd.org/internet/ieconomy/privacy-guidelines.htm`.

In developing the guidelines, the OECD recognized the need to balance commerce and other legitimate activities with privacy safeguards. Further, the OECD recognizes the tremendous change in the privacy landscape with the adoption of data breach laws, increased corporate accountability, and the development of regional or multilateral privacy frameworks.

Asia-Pacific Economic Cooperation Privacy Framework

The Asia-Pacific Economic Cooperation (APEC) Privacy Framework establishes a set of common data privacy principles for the protection of personally identifiable information as it is transferred across borders. The framework leverages much from the OECD Privacy Guidelines but places greater emphasis on the role of electronic commerce and the importance of organizational accountability. In this framework, once an organization collects personal information, the organization remains accountable for the protection of that data regardless of the location of the data or whether the data was transferred to another party.

The APEC Framework also introduces the concept of proportionality to data breach—that the penalties for inappropriate disclosure should be consistent with the demonstrable harm caused by the disclosure. To facilitate enforcement, the APEC Cross-border Privacy Enforcement Arrangement (CPEA) provides mechanisms for information sharing among APEC members and authorities outside APEC.

It's beyond the scope of this book to go into much depth about any of these particular frameworks, legal systems, or regulatory systems. Regardless, it's important that as an SSCP you become aware of the expectations in law and practice, for the communities that your business serves, in regard to protecting the confidentiality of data you hold about individuals you deal with.

PII and NPI

Many information security professionals are too well aware of personally identifiable information (PII) and the needs in ethics and law to protect its privacy. If you've not

worked in the financial services sector, you may not be aware of the much broader category of nonpublished personal information (NPI). The distinction between these two seems simple enough:

- PII is that information that is used to identify, locate, or contact a specific person.
- NPI is all information regarding that person that has not been made public and is not required to be made public.

However, as identity and credential attacks have grown in sophistication, many businesses and government services providers have been forced to expand their use of NPI as part of their additional authentication challenges, when a person tries to initiate a session with them. Your bank, for example, might ask you to confirm or describe some recent transactions against one of your accounts, before they will let a telephone banking consultation session continue. Businesses may issue similar authentication challenges to someone calling in, claiming to be an account representative from a supplier or customer organization.

Three important points about NPI and PII need to be kept in mind:

- **Legal definitions are imprecise and subject to continuous change.** Many different laws, in many jurisdictions, may directly specify what types of information are considered as PII or NPI. Other laws may make broad categorical statements about what is or is not PII or NPI. These laws are updated often and subject to review by the courts in many nations.
- **Doing business in a jurisdiction does not require physical presence there**. If your organization has one customer or supplier in a jurisdiction – possibly even a single prospective such relationship – that government may consider its laws and regulations now apply to you. Ignoring this is a frequent and costly mistake that many businesses make.
- **Persons include companies and organizations as well as natural people.** Businesses and organizations share significant quantities and types of information with each other, much of which they do not wish to have made public. Privacy considerations and the need for information security protections apply here, as well as they do to data about individual people.

It may be safest to treat all data you have about any person you deal with as if it is NPI, unless you can show where it has been made public. You may then need to identify subsets of that NPI, such as health care, education, or PII, as defined by specific laws and regulations, that may need additional protections or may be covered by audit requirements.

Private and Public Places

Part of the concept of privacy is connected to the *reasonable expectation* that other people can see and hear what you are doing, where you are (or are going), and who might be

with you. It's easy to see this in examples: Walking along a sidewalk, you have every reason to think that other people can see you, whether they are out on the sidewalk as well, looking out the windows of their homes, offices, or passing vehicles. The converse is that when out on that *public* sidewalk, out in the open spaces of the town or city, you have no reason to believe that you are *not* visible to others. This helps differentiate between *public places* and *private places*.

- Public places are areas or spaces in which anyone and everyone can see, hear, or notice the presence of other people and observe what they are doing, intentionally or unintentionally. There is little to no degree of control as to who can be in a public place. A city park is a public place.
- Private places are areas or spaces in which, by contrast, you as owner (or person responsible for that space) have every reason to believe that you can control who can enter, participate in activities with you (or just be a bystander), observe what you are doing, or hear what you are saying. You choose to share what you do in a private space with the people you choose to allow into that space with you. In law, this is your reasonable expectation of privacy, because it is "your" space; and the people you allow to share that space with you share in that reasonable expectation of privacy.

Your home or residence is perhaps the prime example of what we assume is a private place. Typically, business locations can be considered private in that the owners or managing directors of the business set policies as to whom they will allow into their place of business. Customers might be allowed into the sales floor of a retail establishment but not into the warehouse or service areas, for example. In a business location, however, it is the business owner (or its managing directors) who have the most compelling reasonable expectation of privacy, in law and in practice. Employees, clients, or visitors cannot expect that what they say or do in that business location (or on its IT systems) is private to them and not "in plain sight" to the business. As an employee, you can reasonably expect that your pockets or lunch bag are private to you, but the emails you write or the phone calls you make while on company premises are not necessarily private to you. This is not clear-cut in law or practice, however; courts and legislatures are still working to clarify this.

The pervasive use of the Internet and the web and the convergence of personal information technologies, communications and entertainment, and computing have blurred these lines. Your smart watch or personal fitness tracker uplinks your location and exercise information to a website, and you've set the parameters of that tracker and your web account to share with other users, even ones you don't know personally. Are you doing your workouts today in a public or private place? Is the data your smart watch collects and uploads public or private data?

"Facebook-friendly" is a phrase we increasingly see in corporate policies and codes of conduct these days. The surfing of one's social media posts, and even one's browsing histories, has become a standard and important element of prescreening procedures for job placement, admission to schools or training programs, or acceptance into government or military service. Such private postings on the public web are also becoming routine elements in employment termination actions. The boundary between "public" and "private" keeps moving, and it moves because of the ways we think about the information, not because of the information technologies themselves.

GDPR and other data protection regulations require business leaders, directors, and owners to make clear to customers and employees what data they collect and what they do with it, which in turn implements the separation of that data into public and private data. As an SSCP, you'll probably not make specific determinations as to whether certain kinds of data are public or private; but you should be familiar with your organization's privacy policies and its procedures for carrying out its data protection responsibilities. Many of the information security measures you will help implement, operate, and maintain are vital to keeping the dividing line between public and private data clear and bright.

Privacy versus Security, or Privacy *and* Security

It is interesting to see how the Global War on Terror has transformed attitudes about privacy throughout the Western world. Prior to the 1990s, most Westerners felt quite strongly about their individual rights to privacy; they looked at government surveillance as intrusive and relied upon legal protections to keep it in check. "That's none of your business" was often the response when a nosy neighbor or an overly zealous official tried to probe too far into what citizens considered as private matters. This agenda changed in 2001 and 2002, as national security communities in the United States and its NATO allies complained bitterly that legal constraints on intelligence gathering, information sharing, and search and seizure hampered their efforts to detect and prevent acts of terrorism. "What have you got to hide," instead, became the common response by citizens when other citizens sought to protect the *idea* of privacy.

It is important to realize several key facets of this new legal regime for the 21st century. Fundamentally, it uses the idea that international organized crime, including the threat of terrorism, is the fundamental threat to the citizens of law-abiding nations. These new legal systems require significant information sharing between nations, their national police and law enforcement agencies, and international agencies such as the OECD and Interpol, while also strengthening the ability of these agencies to shield or keep secret their demands for information. This sea change in international governance started with the Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001, known as the USA PATRIOT Act. This law created the

use of National Security Letters (NSLs) as classified, covert ways to demand information from private businesses. The use of NSLs is overseen by the highly secret Foreign Intelligence Surveillance Court, which had its powers and authorities strengthened by this Act as well. Note that if your boss or a company officer is served with an NSL demanding certain information, they cannot disclose or divulge to *anyone* the fact that they have been served with such a due process demand. International laws regarding disclosure and reporting of financial information, such as bank transactions, invoices and receipts for goods, and property purchases, are also coming under increasing scrutiny by governments.

It's not the purpose of this chapter to frame that debate or argue one way or another about it. It is, however, important that you as an information security specialist within your organization recognize that this debate is not resolved and that many people have strongly held views about it. Those views often clash with legal and regulatory *requirements* and *constraints* regarding monitoring of employee actions in the workplace, the use of company information or information systems by employees (or others), and the need to be responsive to digital discovery requests of any and every kind. Those views and those feelings may translate into actions taken by some end users and managers who are detrimental to the organization, harmful to others, illegal, unethical, or all of these to a degree. Such actions—or the failure to take or effectively perform actions that *are* required—can also compromise the overall information security posture of the organization and are an inherent risk to information security, as well as to the reputation of the organization internally and externally.

Your best defense—and your best strategy for defending your company or your organization—is to do as much as you can to ensure the full measure of CIANA+PS protections, including accountability, for all information and information systems within your areas of responsibilities.

Nonrepudiation

The fundamental design of the earliest internetworking protocols meant that, in many cases, the sender had no concrete proof that the recipient actually received what was sent. Contrast this with postal systems worldwide, which have long used the concept of registered mail to verify to the sender that the recipient or his agent signed for and received the piece of mail on a given date and time. Legal systems have relied for centuries on formally specified ways to *serve process* upon someone. Both of these mechanisms protect the sender's or originator's rights and the recipient's rights: Both parties have a vested interest in not being surprised by claims by the other that something wasn't sent, wasn't done, or wasn't received. This is the basis of the concept of *nonrepudiation*, which is the aspect of a system that prevents a party or user from denying that they took an action, sent a message, or received a message. Nonrepudiation does not say that the recipient *understood* what you sent or that they agreed with it, only that they received it.

NOTE You can think of nonrepudiation as being similar to submitting your income tax return every year or the many other government-required filings that we all must make no matter where we live or do business. Sometimes, the only way we can keep ourselves from harm is by being able to prove that we sent it in on time and that the government received it on time.

Email systems have been notorious for not providing reliable confirmation of delivery and receipt. Every email system has features built into it that allow senders and server administrators to control whether read receipts or delivery confirmations work reliably or correctly. Email threads can easily be edited to show almost anything in terms of sender and recipient information; attachments to emails can be modified as well. In short, off-the-shelf email systems do not provide anything that a court of law or an employment relations tribunal will accept as proof of what an email user claims it is.

Business cannot function that way. The transition from postal delivery of paper to electronic delivery of transactions brought many of the same requirements for non-repudiation into your web-enabled e-business systems. What e-business and e-commerce did *not* do a very good job of was bringing that same need for nonrepudiation to email.

There are a number of commercial products that act as add-ons, extensions, or major enhancements to email systems that provide end-to-end, legally compliant, evidence-grade proof regarding the sending and receiving of email. A number of national postal systems around the world have started to package these systems as their own government-endorsed email version of registered postal mail. Many industry-facing vertical platforms embed these nonrepudiation features into the ways that they handle transaction processing, rendering reams of fax traffic, uncontrollable emails, or even postal mail largely obsolete.

Systems with high degrees of nonrepudiation are in essence systems that are auditable and that are restricted to users who authenticate themselves prior to each use; they also tend to be systems with strong data integrity, privacy, or confidentiality protection built into them. Using these systems improves the organization's bottom line, while enhancing its reputation for trustworthiness and reliability.

Authentication

This element of the classic CIANA set of information security characteristics brings together many threads from all the others regarding *permission to act*. None of the other attributes of information security can be implemented, managed, or controlled if the system cannot unambiguously identify the person or process that is trying to take an action involving that system and any or all of its elements and then limit or control their

actions to an established, restricted set. Note that the word *authentication* is used in two different ways.

- Information is *authenticated* by confirming that all of the metadata about its creation, transmission, and receipt convey that the chain of trust from creator through sender to recipient has not been violated. Authentication of a sent email or file demonstrates that it was created and sent by a known and trusted person or process. This requires that access control as a process grants permission to users or the tasks executing on their behalf to access a system's resources, use them, change them, share them with others, or create new information assets in that system.
- In access control terms, *authentication* validates that the requesting subject (process or user) is who or what they claim that they are and that this identity is known to the system. *Authorization* then allows that authenticated identity to perform a specific set of tasks. Taken together, this is what determines whether you are using someone else's computers or networks with their permission and approval or are trespassing upon their property.

1984 was a watershed year in public law in this regard, for in the Computer Fraud and Abuse Act (CFAA), the U.S. Congress established that entering into the intangible property that is the virtual world inside a computer system, network, or its storage subsystems was an action comparable to entering into a building or onto a piece of land. Entry onto (or into) real, tangible property without permission or authority is criminal trespass. CFAA extended that same concept to unauthorized entry into the virtual worlds of our information systems. Since then, many changes to public law in the United States and a number of other countries have expanded the list of acts considered as crimes, possibly expanding it too much in the eyes of many civil liberties watchdogs. It's important to recognize that almost every computer crime possible has within it a violation of permissions to act or an attempt to fraudulently misrepresent the identity of a person, process, or other information system's element, asset, or component in order to circumvent such restrictions on permitted actions. These authenticity violations are, if you would, the fundamental dishonesty, the lie behind the violation of trust that is at the heart of the crime.

Safety

SUNBURST and other attacks in 2020 and 2021 highlighted how little attention many organizations were paying to the physical control and interaction side of their information systems. Enterprises that did not directly use IT to control manufacturing systems, or command and control vehicles and heavy machinery, believed themselves safe from physical harm, and yet would blithely invest in smart buildings and IoT devices for use in their office environments. They believed that these operational technologies—OT systems that directly cause physical motion or action, or monitor and supervise systems

that do—were sufficiently separated from their IT systems, such as their corporate data centers, that there was little danger of a vulnerability on one side of that IT-OT interface from causing harm to systems, data, and people on the other side. That has been shown to be a false hope.

Operational technologies (OT) include industrial control systems (ICS) and the supervisory, control, and data acquisition (SCADA) systems that direct their activities. OT also includes Internet of Things (IoT) devices, autonomous, mobile machines (from custodial devices to chaotic warehouse forklifts), and robots. Most smart city systems, particularly their mass transit, water and sewer, traffic control, and communications management systems are part of the OT world, as are smart building environmental, power, and security management systems at work and in the home. This list of OT use cases grows every day, and in each case, there are data sharing and collaborative control and supervisory linkages with IT systems at many levels. And in most cases, device control involves switching and detecting AC and DC power and signals as part of controlling physical actuators and sensors.

As older OT systems are being phased out, newer systems tend to be making greater use of the Common Industrial Protocol (CIP). This is a feature-rich set of functions that are used within OT architectures to provide management, real-time control, data acquisition, and safety intervention across an architecture. CIP can operate over IP networks, which allows OT regional control workstations to easily interact with organizational IT systems. OT and IT systems both share common problems, such as the challenges of establishing and maintaining a secure supply chain for software, firmware, and hardware updates. Access control problems are quite common; the information security hygiene measures you need to apply to almost every IT systems environment must also be applied to your organization's OT systems, although with different techniques and tools. Integrated visibility—having a SIEM-like insight into the combined IT / OT architecture of your organization—can be achieved, but it's not as straightforward as some vendors may make it seem.

Safety, like security, is an end-to-end responsibility. It's no wonder that some cultures and languages combine both in a single word. For example, in Spanish *seguridad* unifies both safety and security as one integrated concept, need, and mind-set.

Fundamental Security Control Principles

Several control principles must be taken into account when developing, implementing, and monitoring people-focused information security risk mitigation controls. Of these, the three most important are need to know, separation of duties, and least privilege. These basic principles are applied in different ways and with different control mechanisms. However, a solid understanding of the principles is essential to evaluating a control's effectiveness and applicability to a particular circumstance.

Need to Know

Security classification and categorization should be the linch pin that ties together the organization's information security and risk mitigation efforts. It's what separates the highest-leverage proprietary information from the routine, nearly-public-knowledge facts and figures. Information classification schemes drive three major characteristics of your information security operations and administration.

- **Internal boundaries for information control:** Many business processes have "insider knowledge" needed to inform decisions or exert control over risky, hazardous, or sensitive sequences of actions. These can and should be encapsulated with a layer that hides that inside knowledge by allowing controlled "write-up" of inputs and "write-down" of outputs to the points where they interface with other business processes. These boundaries surround data at higher levels, and the trusted processes that can manipulate or see it, from outer, surrounding layers of processes that perforce operate at lower levels of trust. (It's not a coincidence that that sounds like a *threat surface*.)
- **Standards for trust and confidence:** It's only logical to require higher levels of trustworthiness for the people, processes, and systems that deal with our most vital information than we would need for those that handle low-risk information. In most cases, greater costs are incurred to validate hardware, software, vendors, our supply chain, and our people to higher levels of trust and confidence; as with all risk mitigation decisions, cost-effectiveness should be a decision factor. The information classification standards and guide should directly lead to answering the question of how much trustworthiness is enough.
- **Measures of merit for information security processes:** The level of information classification should dictate how we measure or assess the effectiveness of the security measures put in place to protect it.

Taken together these form a powerful set of functional requirements for the design not just of our information security processes but of our business processes as well! But first, we need to translate these into two *control* or *cybernetic* principles.

Least Privilege

Least privilege as a design and operational principle requires that any given system element (people or software-based) has the minimum level of authority and decision-making capability that the specifically assigned task requires, and no more. This means that designers must strictly limit the access to and control over information, by any subject involved in a process or task, to that minimum set of information that is required for that task and no more. Simply put, least privilege implements and enforces need to know.

A few examples illustrate this principle in action.

- A financial disbursements clerk, when generating payments against invoices from suppliers, has to access and use information about each supplier account as well as access his company's bank-related systems to make the payment take place. However, this clerk would not be expected to modify the information about where the payment should be sent, edit the invoice, or alter the amount of the payment. Nor would this clerk be expected to need any information about other employees, such as their payroll information, while generating payments to suppliers.
- A process control system that actively manages a chemical processing system for a paint manufacturer would not normally be expected to access the Internet or have a need to run web searches of any kind.

Each time you encounter a situation in which a person or systems element is doing something in unexpected ways—or where you would not expect that person or element to be present at all—is a red flag. It suggests that a different role, with the right set of privileges, may be a necessary part of a more secure solution.

Least privilege should drive the design of business logic and business processes, shaping and guiding the assignment (and separation) of duties to individual systems and people who accomplish their allocated portions of those overall processes. Driven by the Business Impact Analysis (BIA), the organization should start with those processes that are of highest potential impact to the organization. These processes are usually the ones associated with achieving the highest-priority goals and objectives, plus any others that are fundamental to the basic survival of the organization and its ability to carry on day-to-day business activities.

Separation of Duties

Separation of duties is intrinsically tied to accountability. By breaking important business processes into separate sequences of tasks and assigning these separate sequences to different people, applications, servers or devices, you effectively isolate these information workers with accountability boundaries. It also prevents any one person (or application) from having end-to-end responsibility and control over a sequence of high-risk or highly vulnerable tasks or processes. This is easiest to see in a small, cash-intensive business such as a catering truck or small restaurant.

- The individual server takes orders, serves the food, and hands the bill to the patron, who pays either the server or the cashier; the cash collected goes into the cash drawer, which is also where any change due the patron is taken from.
- The cash register or change drawer is set up and counted by a shift lead or manager and counted again at intervals throughout the shift and at shift change.
- The daily accounting of cash at start, bills to patrons and receipts paid, and cash drawer tallies is reconciled by the accounting manager.

- The accounting manager prepares the daily cash deposit, which is verified by the overall manager.
- The deposit is counted by the bank and verified to match what is claimed on the deposit slip.

This system protects each worker from errors (deliberate or accidental) made by other workers in this cash flow system. It isolates the error-prone (and tempting!) cash stream from other business transactions, such as accounts payable for inventory, utilities, or payroll. With the bank's knowledge of its customer, it may also offer reasonable assurances that this business is not involved in money laundering activities (although this is never foolproof with cash businesses). Thus, separation of duties separates *hazards* or risks from each other; to the best degree possible it precludes any one person, application, system, or server (and thereby the designers, builders, and operators of those systems) from having both responsibility and control over too many hazardous steps in sequence, let alone end to end.

Other examples demonstrate how separation of duties can look in practice.

- Employees in the finance department can access a financial application and create and change transaction records, which in turn generate audit logs, but those end users cannot access the audit logs. Inversely, a security administrator who can access the audit logs cannot create or change transactions in the financial application itself.
- A developer who writes code for a software release can't also be a tester for that same release or be the one who approves or moves that release into the production environment. If a developer put a flaw into that code maliciously, it will be intentionally allowed to pass. If the flaw was introduced accidentally or through ignorance or poor training, there is a chance the developer will just miss it again in testing. Involving a second person in the testing process allows the organization the opportunity to catch the mistake or malicious flaw.
- An emergency has arisen, and an administrator needs to access a superuser account to perform a sensitive task that is far above their normal level of permissions. Authentication to that account requires a specific hardware token, which must be obtained from the shift lead at the SOC. The SOC lead verifies with the administrator's supervisor and/or verifies that there is an outage, incident ticket, or some other valid reason why the superuser token must be used before issuing it, and the token must be returned when the incident is resolved.

While implementing separation of duties, keep a few additional considerations in mind. First, separation of duties must be well-documented in policies and procedures. To complement this, mechanisms for enforcing the separation must be implemented to match the policies and procedures, including access-level authorizations for each task

and role. Smaller organizations may have difficulty implementing segregation of duties, but the concept should be applied to the extent possible and logistically practical. Finally, remember that in cases where it is difficult to split the performance of a task, consider compensating controls such as audit trails, monitoring, and management supervision.

Another, similar form of process security is *dual control*. In dual control, two personnel are required to coordinate the same action from separate workstations (which may be a few feet or miles apart) to accomplish a specific task. This is not something reserved to the military's release of nuclear weapons; it is inherent in most cash-handling situations and is a vital part of decisions made by risk management boards, loan approval boards, and many other high-risk or safety-critical situations.

A related concept is *two-person integrity*, where no single person is allowed access or control over a location or asset at any time. Entry into restricted areas, such as data centers or network and communications facilities, might be best secured by dual control or two-person integrity processes, requiring an on-shift supervisor to confirm an entry request within a fraction of a minute.

Separation of duties is often seen as something that introduces inefficiencies, especially in small organizations or within what seem to be simple, well-bounded sequences of processes or tasks. As with any safeguard, senior leadership has to set the risk appetite or threshold level and then choose how to apply that to specific risks.

These types of risk mitigation controls are often put in place to administratively enforce a separation of duties design. By not having one person have such end-to-end responsibility or opportunity, you greatly reduce the exposure to fraud or abuse. In sensitive or hazardous material handling and manufacturing or in banking and casino operations, the work areas where such processes are conducted are often called *no lone zones*, because nobody is allowed to be working in them by themselves. In information systems, auditing, and accounting operations, this is sometimes referred to as a *four-eyes* approach for the same reason. Signing and countersigning steps are often part of these processes; these signatures can be pen and ink or electronic, depending upon the overall security needs. The needs for individual process security should dictate whether steps need both people (or all persons) to perform their tasks in the presence of the others (that is, in a no lone zone) or if they can be done in sequence (as in a four-eyes sign/countersign process).

Safety considerations also should dictate the use of no lone zones in work process design and layout. Commercial air transport regulations have long required two pilots to be on the flight deck during flight operations; hospital operating theaters and emergency rooms have teams of care providers working with patients, as much for quality of care as for reliability of care. Peer review, structured walkthroughs, and other processes are all a way to bring multiple eyes and minds to the process of producing high-quality, highly reliable software.

Note that separation of duties does not presume intent: It will protect the innocent from honest mistakes while increasing the likelihood that those with malicious intent will find it harder, if not impossible, to carry out their nefarious plans (whatever they may be). It limits the exposure to loss or damage that the organization would otherwise face if any component in a vital, sensitive process fails to function properly.

Separation of duties, for example, is an important control process to apply to the various event logs, alarm files, and telemetry information produced by all of your network and systems elements.

Of course, the total costs of implementing, operating, and maintaining such controls must be balanced against the potential impacts or losses associated with those risks. Implementing separation of duties can be difficult in small organizations simply because there are not enough people to perform the separate functions. Nevertheless, separation of duties remains an effective internal control to minimize the likelihood of fraud or malfeasance and reduce the potential for damage or loss due to accident or other non-human causes.

✔ Separation of Duties and Least Privilege: It's Not Just About Your People!

In many business settings, the dual concepts of separation of duties and least privilege are seen as people-centric ideas—after all, far too much painful experience has shown that by placing far too much trust *and power* in one person's hands, temptation, coercion, or frustration can lead to great harm to the business. Industrial process control, transportation, and the military, by contrast, have long known that any decision-making component of a workflow or process can and will fail, leading also to a potential for great harm. Separation of duties means that autopilot software should not (one hopes!) control the main electrical power systems and buses of the aircraft; nor should the bid-ask real-time pricing systems of an electric utility company direct the CPUs, people, and actuators that run its nuclear reactors or turbine-powered generators.

Air gaps between critical sets of duties—gaps into which systems designers insert different people who have assessment and decision authority—become a critical element in designing safe and resilient systems.

Access Control and Need-to-Know

As you should expect, these key control principles of need to know, separation of duties, and least privilege also drive the ways in which you should configure and manage identity

management and access control systems, as shown in Chapter 2. Best practices for implementing and managing any IAM system include:

- Create hierarchies of groups of user identities and accounts, with privileges assigned to limit users to the least privileges they require for related tasks and functions.
- Use role-based access control as part of your strategies so that one system or user must explicitly re-authenticate as they change roles to perform more privileged sets of tasks.
- Create nonprivileged user accounts and identities for systems administrators, and others with privileged accounts, and enforce their use for tasks that do not require elevated privileges (such as email or routine web page access).
- Separate groups of user identities and accounts (for people and nonhuman elements of your systems) based on separation of duties.
- Thoroughly examine all installed software, and connections to web or cloud-hosted applications platforms to identify any instances in which apps elevate privileges for nonprivileged users who use such apps or connection. Eliminate such elevation or find ways to explicitly control and restrict it.

Job Rotation and Privilege Creep

Job rotation can be a powerful HR investment strategy that leads to increasing the knowledge and skills of a company's workforce while improving retention of quality personnel, but these are not the concerns of the SSCP. From a security perspective, there are many reasons for creating a job rotation policy. These include reducing risks of both insider and external threats, reducing dependence on a single person (who can become a single point of failure), and increasing resiliency for business continuity and disaster recovery (BCDR) purposes. Banking and investment companies, for example, have used (and have sometimes been required by government regulators or by law) such career-broadening or rotations strategies as part of their loss control and fraud prevention mechanisms.

We cannot overstress the importance of carefully managing what should be the temporary changes in user privileges during such job rotations. Far too often, privilege creep resulting from each job rotation (temporary or permanent) ends up with the user accumulating new sets of privileges with each new task, job, or skills-broadening assignment. Over time, this can lead to an individual having far greater insight into and control over the organization's information assets than should ever be allowed.

In practice, job rotation requires cross-training personnel for various positions and tasks within the organization. This may be within a particular business functional area or discipline, or it might involve a temporary transfer of an employee to other areas

within the company. Some of the personnel in the security office, for example, might all be trained on the various roles in that office (such as log analysis, incident response, security training, or systems testing) as an intra-departmental job rotation and then learn more of the company's human resources or product development business via a career-broadening assignment.

Job rotation helps to mitigate insider threats in several ways. It serves as a deterrent for a potentially malicious insider actually committing fraud. In cases where separation of duties would necessitate collusion, job rotation disrupts opportunities for collusion. In cases where a malicious insider has found a way to mishandle data or abuse their access, job rotation disrupts them from doing long-term damage once they've started. The cross-training aspect of job rotation may also aid the overall security effort by reducing the potential for employees/staff to become dissatisfied and possibly become insider threats; skilled personnel appreciate receiving additional training and challenges of new tasks, and increased training opportunities make those personnel more valuable. Increased morale of skilled personnel reduces costs because of turnover and accentuates loyalty to the organization.

Alternatives to job rotation are forced vacation or leave. The logic here is that if a malicious insider is suppressing alarms, changing or erasing audit logs, or conducting any other activity to cover their tracks or support or assist an attack, this activity should be easier to detect if the suspected insider is suddenly forced to stay away from work. During the period of mandatory vacation, that user's account access should be suspended, and a thorough audit/review of their activity should be performed. This is especially important for those users with privileged access. For example, after the U.S. stock market crash and the collapse of its banking systems in 1929, Congressional action established not only such forced vacations but also frequent bank holidays during which banks suspended customer transaction processing while they performed extensive internal systems integrity checks; both mitigated the risks of fraud, embezzlement, and over-extension by the bank or its staff.

Another goal of job rotation is to keep malicious outsiders from being able to learn about your staff over time and trying to target or manipulate them for information or access. Reducing static patterns in personnel taskings and changing access roles repeatedly reduces the opportunity for external actors to subvert particular employees as targets.

Finally, job rotation also greatly improves the resiliency of an organization, essential in successfully executing BCDR actions. During contingency events or disasters, you must assume that some personnel will not be available/capable of performing particular tasks and functions necessary to maintain the organization's critical processes; having other personnel not normally assigned to those functions but trained on how to perform them is a great benefit and vastly increases the likelihood of BCDR response success.

DOCUMENT, IMPLEMENT, AND MAINTAIN FUNCTIONAL SECURITY CONTROLS

Functional security controls implement the risk mitigation decisions that management and leadership have endorsed. The risk assessment and vulnerabilities assessment tasks have led to these decisions; now it's time to make appropriate cost-effective choices about particular controls, thus *operationalizing* those decisions by providing the tools, techniques, systems elements, and procedural step-by-step that the organization's workforce will need as they go about their day-to-day activities.

The organization has already made decisions about which risks to avoid (by not doing business in particular locations or by abandoning particular business processes); it's also recognized some risks must just be accepted as they are, as an unavoidable but still potential cost of doing business. Chapter 3, "Risk Identification, Monitoring, and Analysis" goes into further depth on how information risks are identified and assessed and how organizational leadership makes both strategic, big-picture risk management decisions, as well as planning for risk mitigation and making the resources available to carry out those plans. Management has also transferred what risks it can to other third parties to deal with. What's left are the risks that you and the rest of your organization's security professionals must deal with. You deal with risk using five basic types of controls: deterrent, preventative, detective, corrective, and compensating. Note that there are no hard and fast boundary lines between these types—a fence around the property both deters and prevents attackers from attempting to cross the fence line, while a network intrusion prevention system both detects and attempts to block (or prevent) intrusions on your networks.

Note that this section focuses on *security* controls, which are of course a subset of the larger problem of risk mitigation. From a security controls perspective, you think about these controls as interfering with a human attacker (or their software and hardware minions) who is carrying out an unauthorized intrusion into your information systems or causing damage or disruption to those systems.

Let's take a closer look at each type of control and then examine common issues involved with their implementation, maintenance, and operational use.

Deterrent Controls

Deterrent controls work to dissuade an attacker from initiating or continuing in their efforts to attack your systems, property, information, or people. Their design, deployment, and use should all raise either the perceived costs or risks to an attacker and the actual costs the attacker could face should they choose to persist. Guard dogs off of the leash, free to range around your property (but within a fence line), are an example of a deterrent that offers painful costs to an attacker, while raising the probability of being forcibly detained and subjected to arrest and prosecution as well.

Deterrent controls should provide a variety of capabilities to the security architect by placing barriers (real and perceived) between potential attackers and the systems they defend.

- Visible, tangible barriers, which an attacker can see, sense, or probe, signal that the target is defended.
- This suggests that the barriers are alarmed and monitored, which increases the possibility of an intrusion being detected.
- The barriers suggest to the attacker that greater assets, time, or effort must be expended for their attack to succeed.
- They also suggest that *more* barriers may be encountered, layer upon layer, should the attacker continue in their attempt.

Note the key concept that to be effective, a deterrent control must be visible, observable, and verifiably present to the prospective intruder. It cannot deter an attacker if the attacker doesn't know that it is there! This directly suggests that you're defending against a known group of attackers and that you have some degree of operational threat intelligence data, which you can use in selecting potentially effective deterrent tactics and techniques.

Simple deterrents can be physical controls, such as fences, locked doors and windows, or landscaping and paving that restricts the movement of vehicles and pedestrians onto a protected property or campus. Exterior lighting, including the use of moving spotlights or floodlights, can also provide a deterrent effect. Most physical controls are passive, in that they do not react to an intrusion attempt; active controls would include guard dogs and security controls, for example.

Physically, the architecture of buildings or workspaces make statements about an organization and the work that is performed there. These statements can also be powerful deterrents to would-be attackers. Think about how many modern embassy compounds (and not just the American ones) around the world have been transformed into little fortresses as they've been blast-hardened, surrounded by impact-resisting barrier walls, and armed military personnel or security guards; entry onto such embassy grounds is restricted and tightly controlled in most cases. High technology companies have also made similar architectural deterrent statements with the ways that they design, build, and operate their physical locations. These are definitely not statements of security through obscurity.

Network systems such as firewalls and intrusion detection and prevention systems can act as powerful deterrents by thwarting an attacker's ability to gain meaningful insight via reconnaissance probes or scans. (It's somewhat unfortunate that the line between NIDS and NIPS as product systems has become quite blurred at this point since both apply filtering rules of varying potency to block or restrict traffic from crossing their point

of protection.) Well-trained, highly aware *people* throughout your organization are also effective deterrents when they smoothly deflect social engineering attack attempts, perhaps by guiding unknown callers through a well-rehearsed script to filter out the innocent prospective customer, client, or job seeker from the whaler-wannabee.

Preventative Controls

Preventative (or prevention) controls provide two forms of protection to keep your systems from harm by reducing the probability of an occurrence of a risk or, when it starts to occur, by containing it in such a way as to limit the spread of its disruption or damage. Securely locked doors and windows prevent an intruder from unlawfully entering your home, unless they want to elevate their risk by breaking through the locks, the windows, or the doors in question. The design of interior walls, doors, and utility spaces restricts the speed with which fire can spread from room to room, while reducing or blocking the spread of smoke and heat. This suggests that security architects should use prevention (like deterrence) in layers.

Prevention can be active or passive, as with deterrence; the same types of controls used for physical, passive deterrence also bring some prevention with them.

Host-based or network-based firewalls, intrusion detection and prevention systems, and of course identity management and access control systems are the main components of a solid prevention architecture. Layer upon layer, they detect attempts to cross a threat boundary's controlled access points; they test that access attempt against varying sets of criteria and in some cases issue challenges requesting further credentials from the requesting subject. Since all of these systems can and should generate both accounting log information for successfully authenticated attempts, and alerts or alarms for failures, they are deterrent, prevention, and detection systems all at the same time.

Detective Controls

Detective (or detection) controls look for any out-of-limits conditions, such as signatures associated with an intrusion attempt, and then take two fundamental and important actions. First, the detection controls notify operations personnel or higher-level supervisory systems that a problem exists; this is absolutely critical if you are to have any command and control over your systems or any ability to manage an effective response to incidents as and when they occur. Second, the detection controls can (if desired) signal an attacker that you've noticed what they're doing, which leads them to believe you'll be responding to their attack. This may deter them from continuing their efforts.

All intrusion or incident detection systems are subject to error rates. Getting the crossover point set so that your risk of harm or loss due to false acceptance errors is balanced by your ongoing costs of investigating and resolving false rejections (and their

concomitant "sky is falling" feeling) is a never-ending process. In fact, the smarter these controls get—and the more that they employ machine learning and predictive analytic capabilities—the more time you'll have to invest in understanding their behavior and tuning it to fit your constantly changing threat landscape and the dynamic nature of your routine business activities.

Physical detection systems can include motion detectors, motion switches on doors and windows, and continuity circuits embedded or built into walls, fences, and other landscaping features. Many such systems can support change detection as well, which can highlight suspicious portions of the systems they surveil to human security monitors for analysis and possible action. Physical systems such as power conditioning, air and environmental conditioning systems, and other aspects of your data center or network operations facilities should be primary sources of alarms that indicate a potential disruption, possibly due to an intrusion, is underway.

Don't forget the end-user element! Properly motivated and trained, having a cadre of end users who can spot something that's not quite right *and* appreciate that management wants to hear about it sooner rather than later can often stymie an attack before it gets too far.

Corrective Controls

Corrective controls provide for the containment, isolation, or restoration of services that have been disrupted for any reason. Uninterruptible power supplies (UPSs) are a good example of this: They isolate or buffer your IT and communications systems from external commercial electrical power providers and in doing so can correct for temporary undervoltage, overvoltage, spikes, noise, or other problems with power before those problems pop circuit breakers or damage equipment. Power problems, incidentally, can also cause equipment to operate in degraded ways that are oftentimes hard to diagnose. Consumer and small business-grade routers, switches, and servers, for example, are prone to odd and intermittent outages for this reason, and the simple expedient of putting them onto an inexpensive battery backup power conditioner or UPS can save hours of fruitless troubleshooting.

Another example of a corrective control in action is when your access control system or a web page design remediates or quarantines a subject's access request when information about that subject and that access request indicates that something is not quite right. Systems can interrogate the subject's endpoint device, for example, to determine whether its operating system, applications, antimalware, or other functions are all properly updated, and if not, route the connection to a remediation server or page that only allows for repair actions to be taken. User subjects can also be challenged to provide further authentication credentials, if something about the time of day, the user's geographic position, or other criteria dictate the need for enhanced vigilance.

Compensating Controls

Compensating controls are put in place when the normal, recommended, or required "best choice" of a risk mitigation control is not available or is unworkable or not affordable or when another approach has been chosen for valid reasons. Depending upon the source of the original requirement for that control, this may or may not be an issue. NIST documents, for example, tend to focus on the risk or threat to protect against, rather than attempting to specify a specific approach. (Best practices, though, often rule out approaches that are no longer useful to consider.) Another example of this can be seen in the Payment Card Industry Data Security Standard (PCI DSS), which specifies stringent security functional or performance standards by which controls must operate, as well as a formalized process for justifying the use of an alternative approach.

PCI DSS gives a good working definition of a compensating control, which can easily apply to other information risk control situations. A compensating control must do the following:

- Meet or exceed the intended level of protection as specified in the original control requirement
- Provide a level of protection that sufficiently offsets or covers the risk that the original control requirement should address
- Must provide greater levels of protection, against the total risk set that the originating or reference standard addresses, than would be achieved by the original control requirement
- Must provide a degree of overall safety and security that is commensurate with the risk of *not* using the recommended or required original standard in whole or in part

This can seem a bit wordy, if not confusing. An example might help. Consider PCI DSS Requirement 3.6.4, as illustrated in a white paper by Robert Schwirtz and Jeff Hall, both at RSM McGladrey. (This paper, which can be found at `https://rsmus.com/pdf/understanding_pci_comp_controls.pdf`, provides good insight into the thinking about compensating controls and how to ensure that a soundly reasoned, well-supported argument is made to justify their use.) This particular requirement specifies that encryption keys must be kept secure. Suppose your system is implemented using a public key cryptography approach such as pretty good privacy (PGP), in which there also is not a centralized certificate authority; there are no keys to keep secure! So, your *compensating* control is the use of a PKI system and the details by which you protect and manage certificates. (Yes, that process involves the use of both parties' private keys, and yes, those have to be kept secure, but these are *not* the keys used to encrypt a PCI DSS transaction. And, yes, it's arguable that the requirement would then apply to keeping the resultant *session keys* secure.)

Another example might be a requirement (in PCI DSS or many other systems requirements specifications) that requires passwords to be of a minimum length and complexity. Using a multifactor authentication system, common sense will tell us, obviates the need for attempts to constrain or dictate user choices of passwords since they are not the sole means of gaining access and privileges.

✔ Residual Risk Isn't "Compensated For"

In common use, we talk about compensating for something as a way to imply that the original would have been better, but for whatever reason, we are settling for less. You compensate for the absence of a key team member by letting others substitute for them, knowing that your team just won't be as strong or the results as good. That's not what *compensating* means when talking about security and risk controls!

For a control to be a compensating control, there is no additional residual risk just because you've replaced the originally required control approach with something different. And if there is a residual risk, then your compensating control is not the right choice.

The Lifecycle of a Control

As with any systems element and the systems themselves, risk mitigation and security controls have a lifecycle that they progress through, from initial observation and expression of a need through implementation, use, and replacement or retirement. More specifically, that lifecycle might include the following:

- Risk identification and characterization
- Vulnerability assessments, with links to specific risks
- Risk management planning decisions, on a per-risk basis, in terms of what to accept, transfer, treat, or avoid
- Risk mitigation decisions, including specifics as to the chosen controls and the anticipated residual risk after the controls are put into practice
- Success criteria, in operational terms, which indicate whether the control is successfully performing its functions
- Anticipated ongoing costs and efforts to use and maintain a set of controls
- End-user and support team training, including any requalification training, needed to keep the controls operating effectively
- Continuous, ongoing monitoring of operational use of the controls

- Ongoing periodic or random assessment, including penetration testing, aimed at assessing the controls
- Decisions to upgrade, replace, or completely retire a set of controls

As you'll see in Chapter 3, there are a number of information products generated by risk management and risk mitigation planning. Although they may be known by various names or be produced in many different formats, the core set of information includes the business impact analysis, risk assessment, risk mitigation plan, and the change management and baseline documentation for the chosen and implemented controls. These could include vendor-supplied manuals as well as your organization's own functional performance requirements allocated to a particular control.

PARTICIPATE IN ASSET MANAGEMENT

Effective information systems management must achieve three distinctly different goals:

- Are we spending what we need to (and no more) to achieve the right business priorities and objectives?
- Are we using our information systems effectively in ways that help us achieve our objectives?
- Are we maintaining, changing, or upgrading our information systems in effective ways to meet changing conditions and needs?

Those three questions all focus on our information systems architecture, the elements we've brought together to create those systems with, and the business logic by which we use those systems. As we'll see in Chapter 3, having a solid baseline that captures and describes our organization's information systems and IT architecture is the foundation of how we manage those information systems. It's also worthwhile to consider that well-managed systems are often more reliable, resilient, safe and secure; unmanaged systems may be just as trustworthy, but if they are, it's more by luck than by design.

Information systems asset management comprises all of the activities to identify each asset, know and control its location and use, and track modifications, changes, or repairs done to it. Asset management also includes keeping track of any damages or losses that an asset incurs through accident, failures of other systems or business functions, misuse, abuse, or attacks of any kind. Due care and due diligence require asset management to be effective, thorough, and accountable, which in turn require that proper inventory and tracking records be kept and that standards be set for proper usage, routine maintenance and repair, safety, and security. Asset management and configuration management and control go hand in hand as the main processes you should use to keep these important,

value-producing assets working well and working for you; they're also crucial to keeping those assets being used by someone else!

ISO 55000 provides extensive guidance for the proper management of physical assets, including buildings, facilities, and infrastructure elements such as electrical power, plumbing, and heating, ventilation, and air conditioning (HVAC) systems. COBIT5, from ISACA (previously known as the Information Systems Audit and Control Association, but now by its acronym only), is another framework of structured guidance for information systems and information asset management, which your organization may already be using.

Broadly speaking, an information systems asset is any element of a system for which it is useful to assess or estimate a value, a cost, and a loss or impact. The value should relate to the gains to the organization that can be realized through effective use of that asset. Costs should reflect all that was spent, including time and effort, to create or acquire, install, use, and maintain the asset. The loss or impact can reflect either the replacement cost, the decrease in value, or some other assessment of how damage, destruction, or degradation of the asset will affect the organization.

Nominally, an asset has one point of management: You manage a single server or you manage a data center, but two data centers integrated via a VPN connection supported by a third party is most likely easier to manage as a set of related assets.

Parts or Assets?

At some point it is easier and more meaningful to track and manage a system as an asset but consider all of the replaceable bits and pieces of it as units or parts. Your network backbone, for example, may consist of high-capacity, redundant routing and switching elements tied together with fiber, cable, WiFi, or other media. As a system, it's useful to track it as an asset, while having a logically distinct inventory of its spare parts.

Asset Inventory

Information systems asset management starts with the asset inventory, which must completely and unambiguously identify every information systems element to be managed as an asset. The inventory should include hardware, firmware, software, virtual machine environments, cloud systems services, databases, websites, and the supporting documentation for end users and maintainers.

Having a current and complete inventory is the absolute bedrock for implementing and monitoring technical security controls.

Robust asset inventory tools and processes will also inform the organization of unauthorized assets. These may be unlicensed copies of software or uncontrolled devices,

software, or systems used by employees, clients, or visitors that thus become parts of your system. They may also be elements of an intrusion in progress. Each of these situations could be risks to the overall safety, security, and reliability of your IT systems.

Note that almost any device that can attempt to access your networks or systems is an object to be inventoried, placed under configuration control, and incorporated into your access control systems' databases as an authenticated identity. Failing to tie these three processes together—and keep them tied together—leaves an unnecessary degree of access open to potential intruders.

Inventory Tool/System of Record

Because of the size, complexity, and frequency of the task, an organization should use automated tools to assist in creating and maintaining the asset inventory. The tools should have awareness of all assets in the organization's enterprise and the ability to discover new assets introduced to the environment that have not been properly documented in the inventory. This data comes from either an asset management agent or a client installed on each asset or "baked in" to each system image. It can also be generated with various scanner and sensor tools, or, in the case of hosted or cloud assets, from a data feed or recurring report from the vendor (which may or may not be shared with clients, depending on the terms of their service-level agreements [SLAs] or terms of reference [TORs] with their clients).

An asset inventory tool should have a way to distinguish authorized devices and applications from unauthorized devices and an ability to send alerts when the latter are discovered. The tool should also collect and track individual asset details necessary for reporting, audits, risk management, and incident management. These details need to cover technical specifications, such as the following:

- Hardware
  - Manufacturer
  - Model number
  - Serial number
  - Physical location
  - Number and type of processors
  - Memory size
  - Network interfaces and their MACs and IPs
  - Hostname
  - Hypervisor, operating systems, containers, virtual images running on this device

- Purchase date, warranty information
- Last update dates (firmware, hypervisor, etc.)
- Asset usage metrics

- Software
  - Publisher
  - Version number, service pack/hotfix number, and date of last update
  - Digital signatures on installation packages
  - License information
  - Purchase date
  - Install date

In addition, operational security details should be collected, such as the type of data stored and processed on the asset, the asset classification and special handling requirements, the business processes or missions it supports, and the owner, administrators, end users, or user groups nominally authorized to use it, and their contact information.

There are of course many tools available that do these tasks or portions of these tasks. Most organizations already own many such tools. Consider the following:

- An Active Directory or Lightweight Directory Access Protocol (LDAP) server can provide a large portion of this information.
- Other integrated identity management and access control systems can provide some of this information and can be especially useful in identifying assets that aren't under management but are attached (or attempting to attach themselves) to your systems.
- Vulnerability scanners, configuration scanners, and network mapping tools can find and provide basic information about all the hosts in the organization's IP ranges.
- Tools that manage/track software licenses can perform a large portion of this task.
- Data loss prevention (DLP) solutions typically have a discovery capability that can serve this purpose.

For gaps in their available tools, organizations can and do compensate with manual efforts, spreadsheets, and scripting to pull and tabulate asset data. Dedicated asset inventory tools usually provide this functionality and preclude the need for manual data pulls and tool integration.

Regardless of the tool or combination of tools used, there should be one the organization deems authoritative and final so that it can be referenced throughout the organization. The information in this tool needs to be definitive. This is the data source to trust

if there is conflict between what other tools are reporting. This should also be the source used for official reports and other data requests, such as part of an audit.

Process Considerations

Let's now look at some inventory management best practices. First, the organization must define the authoritative inventory list or system of record and define the frequency with which the inventory should be refreshed or updated. In addition to the regular interval inventory updates, it is also a good practice to ensure that the inventory management system is updated, and its administrator notified when assets are installed, removed, or updated/changed in a significant way.

This can be accomplished in a different way for environments that make heavy use of virtualized components, including managed cloud service implementations. In these cases, use of automated tools to seek out, tabulate, and provision assets is often preferable; popular tools include Puppet, Chef, and Ansible.

For on-premises assets, it is often helpful to augment the inventory process with the use of geolocation information/geotags or the use of RFID inventory tags. This can increase the speed and accuracy of locating an asset, especially during an incident when time is critical.

Lifecycle (Hardware, Software, and Data)

Although some legacy systems may *seem* to be lasting forever, it's much more common that information systems assets of every kind have a useful economic life span, beyond which it is just not useful or cost-effective to continue to use it and keep it working. Once past that point, the asset should be disposed of safely, so as to terminate exposing the organization to any risks associated with keeping it or failing to care for it. The typical systems development lifecycle model (SDLC) can be applied to hardware, systems software, applications software, and data in all of its many forms; let's look at this from an asset manager's perspective:

- The requirements phase identifies the key functional and physical performance needs that the system should meet and should link these to the organization's mission, goals, and objectives. When any of these change, the asset manager is one of the stakeholders who evaluates whether the asset is at or past its useful economic life.
- During the design phase, the functional requirements are allocated to individual elements of the design; it's worth considering at this point whether these components of the total system should be tracked as assets by themselves versus tracking the system as a whole or as a single asset.

- Development, integration, and acceptance testing quite often conclude with a list of identified discrepancies that must be tracked and managed. In effect, each open discrepancy at the time of systems acceptance is a lien on the overall value of the system (much as a mortgage or mechanic's lien on your home reduces the equity you would realize from selling your home). Tracking those discrepancies is a form of tracking residual risk.
- Operational use presents an opportunity to appraise the value of the system; finding new uses for it increases its value to the organization as an asset, but if users find better, faster ways to do the same jobs instead, this in effect decreases the value of the asset.
- Maintenance and upgrade actions can extend the useful life of the system while adding to its cost. This is also true for ongoing license payments, whether as per-seat or site-wide licenses for software use.
- Retirement and safe disposal, and the costs associated with these, bring this particular asset's lifecycle and its asset management account to a closed state.

Disposal must deal with the issue of data remanence, which refers to information of any kind remaining in the memory, recording surfaces, physical configuration settings, software, firmware, or other forms. This applies to more than just the familiar disks, tapes, and thumb drives; all hardware devices have many different internal nooks and crannies through which live data flows during use. Old-fashioned cathode ray tube (CRT) displays risked having images burned into their display surfaces. Printers have been known to go to the scrap dealer with fragments of previously printed documents, or impressions on their printing drums and ribbons of what they last printed, still legible and visible. Printed documents may need to be shredded or pulped. As a complication, you may end up having to store these retired assets, at a secure location, while awaiting the time (and money) to have a proper zeroization, purge, or destruction of the element to prevent an unauthorized disclosure from happening.

Hardware Inventory

In many work environments, people and whole workgroups can move around within a large facility. People shift from one workstation to another or to larger (or smaller) spaces in another room or another building; some may even move to a different city or country or travel extensively. Hardware inventory needs to know *logically* and *physically* about each device, be it an endpoint, a server, a peripheral such as a printer or scanner or a removable storage device. Assuming for a moment that no MAC address spoofing or alteration is allowed, the identity of an individual device should remain constant; knowing that it's currently attached via a certain IP address and that it is (or is not)

connecting through a VPN is part of knowing *logically* where it is. But…knowing *physically* what desk or tabletop, rack, room, building, or continent it's on (or in) can be problematic. It's prudent to avoid procedurally intensive ways to address this problem, as the German military found out a few years ago. They went from simply allowing their military and civilian staff to just pick up and move their desktop and laptop computers from office to office, as temporary shifts in duties arose, and instituted a work-order process as a way of capturing location information for their asset inventory. This added days of work as each move had to have a form filled in, which was sent to an approvals and dispatch center; then had to have a worker move the equipment; and finally have the form sent back to the user to sign off that the move was now complete. Attribute-based access control (ABAC) may be a smarter solution to such problems, although it may require endpoints that can be trusted to accurately report their physical location without end-user intervention.

I cannot overstress the need to know the physical location for infrastructure elements such as servers, routers, switches, and such, to as detailed a level as possible. Precious time can be wasted during an incident response by having to search for which room, which rack, and which unit or position in the rack is the device that's been sending up alarms (preferably not sending up smoke signals). It's also especially important to note which power distribution panel or circuit breaker box serves each equipment rack or bay and which power conditioning systems feed which distribution panels or breaker boxes.

Software Inventory and Licensing

Software and firmware come in many different forms; almost without question, all of these forms should be under the right combination of configuration control, configuration management, and asset management. Between those three processes, you'll have a very good chance to know that all of your software elements:

- Have been protected from unauthorized changes
- Have had all required changes, patches, and updates correctly applied
- Have had all outstanding discrepancy reports or change requests reviewed and dispositioned by the right set of stakeholders and managers
- Where each element is, physically and logically, how it's being used, and whether or not it is up to date

You'll also know, for each software element, whose intellectual property it is and whether there are license terms associated with that ownership interest. For each license, you'll need to know the detailed terms and conditions that apply and whether they apply to all copies you've installed on any number of devices or to a specific maximum number of devices; the license may also restrict your ability to move an installed copy to another

system. The license might be *seat limited* to a specific number of individual users or *capacity limited* to a maximum number of simultaneous users, maximum number of files or records, or other performance ceilings.

Many modern applications programs (and operating systems) facilitate this by using digital signatures in their installation processes so that each installed and licensed copy has a unique identifier that traces to the license identifier or key. Software license inventory management tools can easily poll systems on your network, find copies of the application in question, and interrogate that installation for its license and identifier information. This can also find unlicensed copies of software, which might be legitimate but have yet to activate and register their licenses or might be bootleg or unauthorized copies being used.

Proper software license management and software inventory management can often save money by eliminating duplicate or overlapping licenses, or by restricting usage of a particular app or platform strictly to where it's needed.

Data Storage

Whether you think of it as *data* or *information*, it is either in use, in motion, or being stored somewhere in the information architectures and systems you are keeping safe and secure. Data can be used by endpoints, servers, or the infrastructure itself. Data is in motion when it is being transferred across networks, communications links, or even to and from a storage device temporarily attached to an endpoint computer or smartphone. Data can be stored – be at rest – in endpoint devices, in removable media, and in storage subsystems that are part of an on-premise network or hosted in a public or hybrid cloud. Chapter 7, "Systems and Application Security," will look in greater depth at security issues relating to data storage in the cloud and within your networks and their servers. What remains is the vexing problem of data storage on paper and on removable storage media and devices, and when those storage media and paper documents are being moved around.

Information Lifecycle

Information has a natural lifecycle, but as with most things in the IT world, there are many different models for this lifecycle, with different emphasis placed on different phases of the data's existence. For example, ISO 27002 defines this cycle with five phases: creation, processing, storage, transmission, and deletion/destruction (see Figure 1.2). Other models, such as those built into many systems management platforms such as SAP, may combine creation and use with processing; then add a *retention* phase in which the data is not actively used but cannot be disposed of because of legal, regulatory, or liability reasons; and finally end with a disposal and destruction activity.

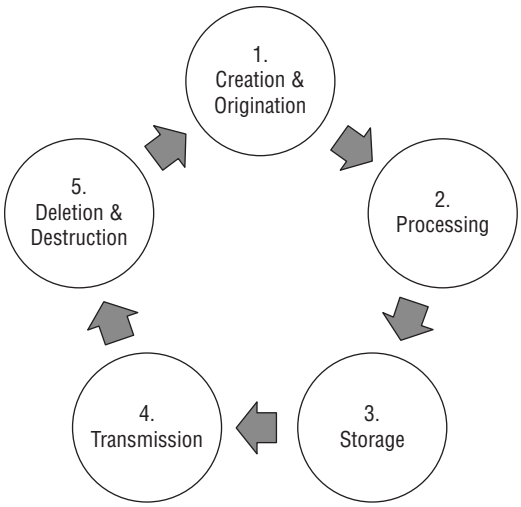

FIGURE 1.2 **ISO 27002 phases**

Security is an important consideration at every phase, but the level of importance can vary, depending on the phase. The formats and media used in the various phases can also affect the security considerations.

Consider, for example, design documents for a new product or technology. When those documents/data are new, they are valuable and actionable, especially if a competitor acquires them. Once the product or technology is in production and is being sold on the market, those same documents could be near the end of their lifecycle. At this point, one could argue that the documents would do less damage in the hands of a competitor, but they still need to be afforded some level of protection, right up to the moment they are destroyed. In this example, even though the creators have benefited from the "rush to market" advantage, the design documents themselves could still contain sensitive data, such as a proprietary production method, which the organization plans to reuse in future products.

There are several important points to take from this example. First, the security impact may vary depending on where the information is in its lifecycle. Next, even though the impact may vary, there can be negative outcomes for the organization at *any* phase. Finally, phase five, the deletion and destruction phase, is important because destruction of unneeded assets reduces the organization's attack surface. Data at the end of its lifecycle only introduces risk, with little to no benefit.

The lifecycle view shows that datasets (or information assets) are constantly going back and forth from storage to use; throughout *that* ever-repeating cycle, your systems designs should protect the information while it is at rest, in use, and in motion. Currently, well-chosen encryption systems can protect data in motion and at rest and by means of digital signatures offer the stored copies protection in time as well. (Chapter 5

goes into this in further detail.) However, thus far there are not many solutions to protect data while it is being used from compromise or loss, since most operations on data and its use by humans needs to have the meaning of the data readily available.

Apply Resource Protection Techniques to Media

Protecting the information on storage media requires that you can control or limit the onward use, copying, or other redistribution of that information; it also requires you to protect your systems from being contaminated by information from a classification level that does not belong on your systems. For example, the Biba and Bell–LaPadula access control models to show how different models emphasize confidentiality or integrity. Both choices can be undone by putting the wrong level of information onto the wrong removable media and then introducing that media into another system. You'll see a variety of standards and practices in use that may place different emphasis on protecting either the information (and its confidentiality, nonrepudiability, or integrity) or the systems (by protecting their integrity, and hence their availability and authenticity).

Before covering the methods for properly managing media, it's important to acknowledge that these methods will vary based on the types of media used. The umbrella term of *media* or *information system media* could mean legacy analog formats, such as hard-copy documents, photos, and microfilm. It could also (more likely) be in reference to a wide range of digital formats, such as external hard drives, floppy disks, diskettes, magnetic tape, memory cards, flash drives, and optical disks such as CDs, DVDs and Blu-Ray disks.

As you might expect, making secure but removable media work requires successfully integrating your security classification schema, your device-level identity management and access control, and the management of all endpoints' capabilities to use removable storage—including on the endpoint itself.

Marking

Handling sensitive or classified information involves everything necessary to meet its protection requirements; handling storage media refers to all processes, be they human, electronic, or mechanical, which are involved in mounting, dismounting, storing, shipping, using, reusing, and ultimately destroying the media. This protection requires a combination of marking the media and establishing and using administrative and logical processes that perform those tasks in controlled, reliable, and auditable ways. The marking achieves nothing without the procedures being understood and used properly!

Marking involves labeling in both human-readable and machine-readable manners so that it is immediately obvious what the highest security classification level of data on that media can be and should be. Humans are known to put "for unclassified use only" disks into drives and then write secret, proprietary, or private data to them, either deliberately as part of an exfiltration attempt or accidentally. The labeling should clearly link to the

proper handling procedures for that level of security classification. Done properly, your device-level identity management and access control systems can then use this marking to authenticate the media when it is first mounted and then authorize each attempt to read or write data or metadata to it.

It's strongly recommended that your IT or security teams be the ones who purchase, label, initialize, and inventory storage media used for sensitive, proprietary or other data security classifications your company uses. When teamed with user-facing policy directives, this can significantly reduce the compromise of classified information due to a user forgetting to properly label a piece of media.

✔ Colorize to Classify

Marking media might become complicated, depending on the media used. For instance, it might be possible to include a significant amount of information on the label of a 3.5" floppy disk, but much, much more difficult to put that same information on the label of a USB flash drive that is the size of a thumbnail. Quite often, it's much more effective to use color schemes as a visible part of media security marking, when the media itself can be readily purchased in a range of colors suitable for your organization's security labeling needs. Many media vendors can also prelabel the physical media itself to meet your needs.

Protecting

Consistent with the least privilege and separation of duties concepts discussed previously, your organization should restrict access to and usage of removable media to specifically authorized staff members who need it for their daily duties, based on their specific roles.

To do this, there must be an element of physical protection and storage that is commensurate with the sensitivity and classification of the data on the media. Here are a few examples, illustrating different levels of protection:

- Backup copies of audit logs are kept in a locked desk drawer or cabinet, where the key is available only to administrators who may need to review the logs.
- Signed hard-copy health insurance forms are in a locked file cabinet in a room restricted to HR staff via proximity-badge access.
- An external hard drive with classified data on it is fully encrypted and is in a locked safe in a protected area, accessible only to users with appropriate security clearance and need to know. The encrypted files can be decrypted only on systems that are cleared for using information at that level and then only when being used by a user with matching privileges.

As you can see in the examples, different layers of both physical and logical access control can and should be provided to media to meet your information security needs. There are additional measures to consider, based on the sensitivity and criticality of your media. You may need to create redundant copies of critical media to mitigate accidental damage or loss. Suitable encryption and other techniques can protect the classified data while it is at rest (stored on the media) and in motion between the media and the systems that are processing it (and making it available to users). Remember, too, that all storage media and technologies suffer degradation over time, resulting in data loss. Your data integrity, availability, and retention needs may drive you to establish a media rotation strategy, which periodically moves the files (the in-use set and the backup copies) to new, fresh media. (Data centers have been doing this since the 1960s, as they discovered that reels of magnetic tape quite literally saw bits flaking off when they hung in storage for too long.) Finally, you should treat the collection of all of your sensitive, critical information and the media it is stored on as a library of assets and define formal processes for periodically verifying your inventory of media, for formally authorizing users to check media in and out of the media library, and for leaving an audit trail. These processes should be followed until the media is either sanitized and then downgraded for uncontrolled use (not recommended—it's a false economy!) or destroyed for disposal, using approved equipment and methods in either case.

Transport

Your organization needs to have a defined set of procedures for protecting media when it is transported outside of controlled or restricted areas. These procedures should define the check-in and checkout accountability mechanisms used for transport, as well as the documentation requirements of the transportation activities. You should also explicitly define what information must be captured or logged upon checkout, during transport, and upon check-in of media, which might include details such as who requested the transport and who was responsible for the media during transport.

Any staff or courier transporting media should clearly understand the restrictions applied to the transport (such as approved travel methods, routes) as well as special handling and packaging considerations, based on media type, to protect it from hazards such as moisture, temperature, and magnetic fields. This also includes when, whether, and how encryption should be used during transport. Couriers should also understand your rules on deviations from procedures in the event of unforeseen circumstances encountered during such transport.

Transport procedures should be clear as to when appointed custodians are necessary, who the approved custodians or couriers are, and how to verify identity if external couriers are used. Consideration should also be given to when and how the responsibilities of the custodian can be transferred to another, as well as specific points of contact to whom the media can be transferred at arrival.

Sanitization and Disposal

The topics of media sanitization and disposal overlap and are interrelated. There is a time in the information lifecycle when certain data is no longer needed, and having this data sitting on media for no reason presents an unacceptable risk. If there is no benefit, why accept even the slightest risk that the media could be compromised? At that point, the information must be destroyed by sanitizing or zeroizing the media; the media may be returned to your library as reformatted, empty, but suitable for reuse with information at a security level consistent with the media's marking or destroyed if the media is past its economically useful life as well. So, what are the differences between the two?

The first difference is the reuse scenario. According to NIST 800-53, media should be sanitized "prior to disposal, release out of organizational control, or release for reuse." Disposal of media doesn't acknowledge a need to reuse the media, but sanitization does. Blank, new media might cost $50 to $3,000 or more apiece, so it may be worthwhile to have effective reuse and sanitization strategies in place. With the rapidly increasing capacity and decreasing cost of solid-state drives and flash media, many organizations choose verifiable destruction rather than risk an incomplete sanitization of such media. Destruction can also be done faster and at less cost in most cases.

The next difference is in the methods. The sanitization methods are less physically destructive than disposal methods. For example, sanitizing nondigital media, such as paper documents, is accomplished by removing sensitive pages or entire sections or by redacting or obscuring specific text. In contrast, disposal of paper documents would entail cross-shredding, pulping, or burning the papers entirely. Sanitizing digital media, such as hard drives, would mean overwriting each sector and byte of the drive many times with random characters. (The NSA has been known to call this process *zeroization*, even though it doesn't actually recommend writing nothing but zeros to the media; this would risk a missed block or sector being completely readable.) Disposal of hard drives, in contrast, entails either degaussing the drive, physically abrading or chemically corroding the surface of the disk platters, or breaking the entire drive in a powerful shredder. Even when degaussed or abraded, disposal of sanitized media may be constrained by local laws, including any limitations on the search of trash disposal sites with or without a search warrant.

NOTE Degaussing does not work on a solid-state drive (SSD) or optical disk.

Another slight difference you can see in the NIST verbiage is that sanitization is often a defense-in-depth approach to precede disposal and augment it as a security control. Imagine, for example, a scenario where a hard drive was not effectively destroyed by the organization's normal disposal method or was, for example, intercepted by a curious or

malicious person in the chain of custody. Even if the drive wasn't destroyed but had been previously overwritten many times with random characters, it may still be unreadable, and the sanitization is a good mitigation for the failure in the disposal process.

Having discussed the differences, what are the commonalities between sanitization and disposal? Essentially, everything else. The goal of both sanitization and disposal is to ensure that the data previously on the media is not readable or recoverable. They should both happen according to formal processes that review, approve, document, and verify the sanitization/disposal. In both cases, the methods and tools should be commensurate with the data stored on the media. This also includes the removal of external markings and labels.

For both sanitization and disposal, the sensitivity of the data on the media should drive how rigorously you apply these processes and how thoroughly you control it procedurally. In some cases, also consider that it may be less expensive to apply the more stringent sanitization or disposal method to all media than to spend time separating them.

Both sanitization and disposal use specific tools, whether software tools, grinder, shredder, degausser, etc. These tools need to be periodically tested to ensure they are effective and that the media/remnants cannot be read or restored.

When storing and collecting media *prior* to sanitization or disposal, consider affording additional protection above and beyond normal media classification and marking. If there is a large quantity of nonsensitive information in one place, it can become more sensitive by aggregation.

✔ Media Disposal and Information Retention Must Match

Almost every category of corporate or private-sector sensitive or classified information has to have a *retention strategy* defined for it, as part of keeping the organization compliant with a growing and sometimes bewildering body of law and potentially conflicting stakeholder interests. Make sure that your information library procedures, including the ones for destruction of information and disposal of media, match with those retention requirements. If they don't, you'll need help from senior management and the organization's legal team to find an acceptable solution.

IMPLEMENT SECURITY CONTROLS AND ASSESS COMPLIANCE

Although it seems a bit of an oversimplification to do so, you can characterize the world of information security controls (also known as *risk mitigation controls*) by their mix of physical, technical (or logical), and administrative elements. For example, a perimeter fence

is both a physical investment in a control technology and its accompanying procedures for a periodic inspection, including "walking the fence line" by the security patrols and repairing damage by Mother Nature, vandals, or intrusion attempts. Technical or logical controls are the software and data settings, the jumper plugs or control switches, or other device or system configuration features that administrators use to get the software and hardware to implement a security control decision. Windows-based systems, for example, use software-defined data structures called *group policy objects* (GPOs) that apply logical rules to subjects and objects in the system to exert security control over their behavior. Most network devices are *logically* configured by interacting with their GUI, a built-in web page, or a command-line interpreter, to accomplish the technical configuration of that device so that it does its part in carrying out the organization's security policies.

NOTE It's helpful to remember that a *physical* control interacts *physically* with the subject or object being controlled; technical and logical controls interact with data flows and signals being sent around the system as ways to control the logical behavior of software and hardware.

Chapter 3 will focus on how you choose what mix of physical, logical, and administrative controls to build into your security architecture; here, we'll focus on them after you've installed them and declared them operational.

Regardless of the type of control elements involved, compliance can be measured or assessed by the same set of techniques: review, audit, exercise, and operational evaluation. Help-desk trouble tickets, user complaints or suggestions, the "police blotter" or daily logs kept by your security teams, and many other sources of information should all be subject to review and audit. Performance metrics can also be adopted (preferably in automated ways) that can alert management when controls are not being used effectively, as indicated by increasing rates of incidents, error rates, problem reports, and end-user dissatisfaction with system usability and reliability. Don't forget to keep an eye on customer or client behavior and input: A decline in orders, transactions, or web page hits may be as much about the quality and price of your products as it is about the security (or lack thereof) of your information systems and practices, as seen by your customers.

Technical Controls

In all cases, you should first have an administrative (people-facing) control document, such as a policy statement or a procedure, that provides the justification and the details you need to configure, operate, inspect, and update all of the technical settings that implement the electronic aspects of your security architecture. These include both the networks, servers, and endpoint technologies, such as software Group Policy Objects, parameter files, access control lists, or even jumper and patch panel settings. Also included are the

programming, options, and controls for fire and safety alarm systems, motion detectors, entryway alarms, power conditioning, and environmental control systems. (Remember, it was through a maintenance back door in the heating, ventilation, and air conditioning systems that attackers were able to gain entry into Target's systems in 2013.)

Two of the most common technical controls used in many security strategies are related to setting time limits on user activity. *Session timeouts* or *inactivity lockouts* can be implemented on an endpoint device level, on a per-user ID level, or by individual applications platforms, servers, or systems. They force a user to take a positive action to reconnect or renew a login (and go through some or all authentication factor steps) to continue, once their device, session, or use of that resource has been inactive for a specified period of time. This can be frustrating to users when they've come into a system via SSO authentication, gaining initial access to a set of resources and applications at the start of their workday but then having to repeatedly log back in again when they've let individual sessions with specific apps or servers go idle for too long. Session timeouts provide protection against the "lunchtime attack," which got its name from an intruder being able to wander around an office building and find computers unattended but still logged into the system during lunch breaks. Device-level timeouts on company-managed endpoints are typically set for short periods, such as 10 minutes, based on similar reasoning. Specific applications platforms and the portals that your users access them through may need to impose their own timeout periods and choose whether to use timeout warning reminders, based on specific systems security needs.

Another time-based technical control, the merits of which are hotly debated, is *password aging*; this sets a time period (usually measured in days, not minutes) after which a user must change their password. Other password policy settings can limit password reuse as well. Password aging, length, complexity, or other password characteristics should be determined as part of your integrated approach to identity management and access control; proper implementation of multifactor authentication, for example, may provide greater security and ease of use than complex, rapidly aging passwords were once thought to provide.

All of these settings should be subject to formal configuration management and control and documented in some fashion so that an incident response team, network operations staff, or the IT team can quickly refer to them to determine whether the alarms are sounding due to a misconfigured control or because a security incident is occurring.

Physical Controls

Physical controls are things you can touch (or bump into); they are the walls, doors, locks, fences, mantraps, concrete barriers, and their relative placement in your overall physical arrangement of a facility. By themselves, physical security features provide

deterrent, prevention, and containment capabilities; to get your money's worth out of them, most organizations add monitoring equipment such as cameras, motion detectors, alarms, and people (and perhaps security canine patrols). As more robots and autonomous mobile devices enter the workplace, physical access controls must be able to cope with their presence and movements. Gluing that all together requires administrative controls in the form of policies, procedures, and control documentation. It also relies upon the human element—the monitors, the watch-standers, and the administrative and technical people who make it work and who use it to secure and protect the organization, its people, its information, and its assets.

✔ Human Vigilance—Keep It Working *for* You

Whether you consider it part of your physical or administrative control systems, the human element in your security architecture can and should provide significant return on your investment in it, which you can achieve by treating them as professionals. Recruit them as if they matter to you (which they do!). Make sure that initial onboarding and training informs, empowers, and inspires them.

You have a leadership opportunity with everyone involved with security operations, whether you're their supervisor or not. Step up to that challenge, work with them, and lead them as a team to be part of what keeps *everybody's* jobs secure. No matter what functions they perform or whether they stand around-the-clock watches and patrols or only work normal business hours, they can be pivotal to keeping your systems and your company safe—or become some of the weakest links in your chain of security if you ignore them or let others in the organization treat them shabbily.

And if it's your first day on the job, be sure to treat each and every one of them as the helpful, dedicated professional that they are. The paybacks of this strategy can be unlimited.

Physical security architectures usually place high-value assets and systems within multiple, concentric rings of physical perimeters. Entry onto the property might require going past a guard post; checkpoints at the entries to individual buildings on the property would authenticate the individuals attempting to enter and possibly conduct a search of the personal property such as briefcases or backpacks under their control. (Most jurisdictions do consider that owners or managers of private property have the legal right to require that visitors or staff voluntarily allow a search of their person and belongings and deny entry to those who decline to cooperate with such a search.) Once inside, lateral movement within an area or access to high-value areas such as documentation or software libraries, financial operations centers, server and network rooms, or security operations control

centers are further restricted, perhaps requiring two-person control as part of authentication procedures. Layer by layer, these cascades of control points *buy time* for the defenders, time in which any errors in authentication can be detected or subsequent attempts by the subject to exceed authorized privileges generate alarm conditions.

Controlled entry systems, such as mantraps and turnstiles, are electromechanical systems at heart. On the one hand, these must interface with some portion of your identity management and access control systems to be effective; on the other hand, they need routine maintenance, as well as remedial maintenance when they fail during use. In most cases, human guards or controllers are present in the immediate vicinity of such control points.

Controlled egress systems may employ the same physical, logical, and administrative tools as used to control entry into and movement within a facility; they bring the added benefit of controlling inventory, equipment, software, or data loss (sometimes called shrinkage by wholesale and retail businesses), by both deterring and preventing unauthorized removals from occurring. This usually requires a degree of search of property as it leaves the controlled area. A growing number of high-technology firms, especially in biotechnology, rigorously enforce controlled egress and search as vital components of protecting their intellectual property and competitive advantage.

Video and audio monitoring systems have become standard elements in most security systems—and all the more so as the costs of fully digital systems have become much more affordable. Even the small office/home office (SOHO) entrepreneur can afford a multicamera, digital video recorder security system, complete with Internet interfaces for remote monitoring. Many security cameras now come with infrared LEDs that provide surreptitious illumination of the scene, which improves monitoring significantly without needing to add visible light floodlighting systems and their power distribution and control elements; note that after keeping the lenses clean, proper lighting is essential for useful image quality.

Inspection and maintenance of physical control systems is vital to continued security. Administratively, there should be no surprises here; if a maintainer or inspector shows up, your on-shift, on-site guards and monitors and the security control force all need to first authenticate their identity and further confirm that they've been properly called out or dispatched to perform a specified set of tasks.

All physical control systems elements should be documented and under formal configuration management and control appropriate to their physical nature. Concrete block exterior walls, for example, should not be subject to having holes drilled or cut into them without proper authorization. The security department might not control or manage all of this documentation or the change management processes for the structural elements of the physical security aspects of your systems; regardless, your organization's security needs suggest how closely the building maintenance teams and the security teams need to work with each other.

Administrative Controls

In most organizations and the cultures they are rooted in, there is a natural hierarchy of guidance and direction, starting with broad, sweeping, and visionary statements that get progressively less motivational as they become more prescriptive. Subsequent layers become *proscriptive*, tending to have as many "thou shalt nots" as they have "shall" statements in them (if not more). Although the names for many of these layers may be different in different settings and cultures, it's still reasonably useful to expect the same basic layers of policies, standards, procedures, baselines, and guidelines.

Policies

Policies are at the heart of what the organization is trying to accomplish. At a high level, policies provide critical instruction to senior executive management to implement measures to achieve external compliance expectations or support the larger strategic vision of the organization. This layer of senior management then promulgates these vision statements down to more tactical and operational managers both as policy statements and in finer-grained direction. As governance documents, the responsibility for creating and maintaining policy rests with the board of directors or other formalized group of senior stakeholders and leaders. As such, policies are one of the ways in which the board demonstrates due care. Boards can and often do delegate or direct that executive or operational management develop these policies and bring them back to the board for review and endorsement.

Policies, relative to other organizational documents, are less likely to change. They provide consistency to the organization's management, allowing the leadership to shape standards and create procedures that achieve the policy end. They should provide management with sufficient flexibility to adapt to new circumstances or technologies without a policy revision.

Mature organizations routinely review their policies within their governance processes. Changing external compliance expectations or shifts in business strategy almost always require changes in statements of policy and vision. Additionally, these same external factors may cause the organization to confront or consider changes to their previously established strategic goals and objectives, which will probably drive more policy changes. The policy review process must address the changing needs of external stakeholders to support predictability in execution of the policies by management.

The use of the term *policy* when implementing security practice in an organization is often confusing. For example, a password policy may, or may not, be of interest to the governing organization—but it certainly would be of interest to the management team! The organization's governance structure would likely express interest in ensuring access controls are present and that the compliance expectations are appropriate to the

organization's needs at the policy level and leave to management the decision of how many times a password should be rotated. That management chooses to refer to the outcome of their due diligence as a policy is an organizational decision.

Sometimes referred to as *subpolicies*, these amplifying instructions further set behavior expectations for the organization. Some of the areas that might be addressed include passwords, cryptography, identity management, access control, and a wide range of other topics. The critical distinction is whether the instruction comes from the governance body (making it a policy) or whether it is derived from a higher-level policy by the organization's management.

This broad use of the term *policy* reflects one of the major challenges in our industry. A lack of a common language for information security practice has been repeatedly identified as one of the factors inhibiting the development of a common body of practice in the information security community. It is further complicated in an international environment where translations and cultural differences affect how people perceive information. In addition, the various standards bodies have published specific definitions for information security terms that may have nuanced differences between each other.

And if that's not confusing enough, there are many instances of operating systems configuration settings that are also called *policies*.

Standards

Once the organization has decided what it wants to accomplish, management can start to perform tactical planning and operational activities to carry out the intent of the policies. One tool to support efficient management of resources is the use of standards. Standards simplify management by providing consistency in control. External standards are ones developed outside of the organization, usually by governments or industry association standards-setting bodies such as the IETF or IEEE. These provide the world with a uniform vision, purpose, and set of details about the issues that the standard focuses on. Companies can also generate their own internal standards, which they may choose to make as mandatory on all of their systems. Regardless of where the standards come from, they are downward-directed by management onto lower levels of management and supervision to support the achievement of the organization's strategic goals and are tied directly to the organization's policies. Standards also represent a consensus of best practice, as understood by the body that issues the standard. Standards may also be required as part of legal or regulatory needs or because a contract with a key customer requires the standard to be applied to work performed under that contract.

Private organizations may be required to adopt certain standards to do business in a particular market. For example, if an organization wants a web presence, it has to take into account the standards of the World Wide Web Consortium (W3C) in developing applications.

While standards are a management tool, standards often evolve out of organizational practice. For example, selecting a particular vendor to provide a product may force a standard where none was originally contemplated. De facto standards often evolve inside organizations as different parts of the organization adopt a new technology, not as a conscious management decision.

Well-structured standards provide mechanisms for adaptation to meet local conditions. Through the use of baselines, an organization can shape a standard to better reflect different circumstances. Baselines enable the delegation of decision-making within strict parameters to lower levels of management.

Nevertheless, standards are directive in nature; compliance is not optional. At most, the standard itself and the contractual or legal requirement to abide by it may specify ways in which the application of the standard can be tailored to the task at hand. Organizations that adopt standards may also be required by those standards, by contracts, or by other compliance needs to monitor the successful application of and compliance with those standards.

Procedures

Procedural documents provide highly detailed task-oriented instructions. Procedural documents are useful when a high degree of compliance is necessary and the precise steps to achieve the outcome are not readily apparent to individuals not familiar with the environment.

Management, as part of its diligence responsibilities, enforces organizational procedures through routine oversight and audit. Compliance is not optional, and well-structured organizations track compliance with procedural steps.

In certain environments, procedural compliance is achieved by using various separation-of-duties methods. For example, in cloud environments, an organization might require that every action applied to the cloud environment is performed by using an approved configuration management script, such as a Chef recipe or a Puppet task, while further dictating that the author of a script cannot be the same individual who approves the script.

Note, too, that the word *procedure* is also used by software developers and programming languages to refer to a unit of software, such as a function, a subroutine, or a stored query.

Baselines

Some organizational cultures refer to a tailored version of a standard as a *baseline*. Typically, tailoring of a standard reduces the requirements set by the standard; if additional requirements are needed, it is best practice to put them into some other document, such as a local or internal standard. Once a baseline has been established, any deviation from the baseline should be formally approved through the organization's change management practice. As with standards, baselines establish a compliance expectation.

As a subset of baselines, *security baselines* express the minimum set of security controls necessary to safeguard the information security requirements and 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.

The term *baseline* can also refer to a reference set of systems components; the inventory of software installed on a server by the vendor, at the time when the server is first turned on and configured, is an *architectural* baseline.

Guidelines

Guidelines are necessary when an organization determines that some level of flexibility in implementation is necessary to achieve business objectives. Guidelines often rely upon best practices for a particular discipline or are the codification of an organization's experience in a particular area.

Guidelines may be useful when a range of options exist to achieve a particular control objective and it is acceptable to encourage creativity and to experiment to compare the effectiveness of different options. Guidelines may also be useful when the organization's staff has a broad base of experience and a shared vision for an outcome. In that case, the explicit directions of procedures, standards, and baselines may provide too much structure and impede the adoption of more efficient methods.

There are many sources of guidelines for information security practice. Certainly, the CISSP Body of Knowledge 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.

There is no single, correct answer for the number and breadth of policies, standards, baselines, procedures, and guidelines an organization should have. Different regulatory environments, management expectations, and technology challenges will affect how the organization expresses and achieves its goals.

Periodic Audit and Review

There are two major shortcomings with most human-facing procedural and administrative controls for security and risk mitigation. The first is that in their human-facing form as an end product, they invariably end up being anywhere *but* right at the point of

contact between the humans involved and the vulnerable system element the administrative controls are designed to protect. Policies and procedures distributed on paper or as email attachments end up being lost or buried in a desk drawer or folder tree and forgotten about. Signs and warning placards catch the eye during the first few days or weeks after they've been posted, but after a while, the human mind tunes them out; they're just part of the visual clutter of the background.

Because of these shortcomings, it's good to audit your administrative controls with an eye to separating them into two major categories: those that direct or require a real-time action, such as emergency notification and incident response; and those that provide longer-term guidance for behavior, such as inappropriate or unauthorized use of company-provided assets and resources. That first category represents opportunities for some smart investment to ensure that just the right amount of policy guidance, direction, and constraint is at the right fingertips at the right time.

Audits

Audits are structured reviews that compare a set of security and risk controls, and the systems that they protect, against a controlled administrative baseline. This baseline can include inventories, performance standards, compliance standards and requirements, quality measurements and standards, or process maturity models and standards. Informal audits can be used as part of troubleshooting, to improve organizational knowledge of its own systems, or to gain insight into opportunities for improvement. Informal audits do not require the use of outside auditors who are trained and certified for the type of audit being performed. Formal audits, by contrast, are typically conducted to meet legal, regulatory, or contractual compliance needs, such as those imposed by governments or the organization's finance or insurance providers. Audits produce a report, which is typically addressed to the management or leadership levels of the organization that requested the audit. Although the structure of these reports can vary considerably, they usually include an executive summary of the audit, key findings, issues or discrepancies that need to be resolved, and any recommendations as appropriate.

Audits can place a significant burden on information security operations and support teams. Typically, extensive preparation is required to identify the audit baseline or standards that will be used and ensure that the auditors will be able to access all of the items being audited. Workspaces will need to be provided for the audit team, and the auditors may require special access and privileges to the IT elements being audited. They may also need to have IT systems to use for gathering and organizing audit data and to produce and report their findings.

Exercises and Operational Evaluations

Things change; that is the only constant we have in life. The proficiency and currency of the tacit knowledge within your team changes with time; the threats change how they seek opportunities that meet their needs and how they attempt to exploit them. Your systems change, and sometimes not for the better as they age in place. For these and many other reasons, it's wise to establish a process of exercising and evaluating security and risk mitigation control systems, in as realistic an operational setting as you can manage without unduly disrupting normal business operations. A properly designed and well-considered exercise and operational evaluation plan should gain the support of management and leadership; their guidance and sponsorship are crucial to make time and talent available to plan and conduct such activities. Be sure that each plan closes with a thorough post-event debrief and analysis, producing documented recommendations or action items to finish the job of learning what each exercise or evaluation just finished teaching you and the evaluation team.

PARTICIPATE IN CHANGE MANAGEMENT

Change Management or Configuration Management?

These two terms are quite often confused with each other or used as if they are interchangeable; in point of fact, it depends upon the culture and environment you're in as to which name is best to use. In business and leadership development contexts, change management (and change leadership) involves motivating, guiding, and leading people to change the ways they perceive their work and their systems and then further leading and guiding them toward making changes in those systems and in themselves. In these same contexts, configuration management is about taking a defined set of hardware, software, information, and even people skills and tasks, each of which has its particular collection or configuration of settings, options, parameters, and feature selections, and changing it into the same set of elements with different configuration settings. When you talk about IT change management and what you really mean is changing an IT systems' technical configuration into another configuration, it may be less confusing to talk about this as IT configuration management rather than IT change management. (Fortunately, nobody seems to talk about leading people to behave differently as "reconfiguring" them or managing that growth and development as "configuration managing" the HR assets!)

In an effort to reduce confusion, throughout this book I will refer to decisions about changing the configuration settings of an IT system as *configuration management*. (Change management, in the sense of organizational mission, vision, purpose, and culture, is beyond the scope of this book.)

As with many other topic areas, configuration and change planning and management present opportunities for you to work with the people around you, and with the procedures they already have in place, to understand what meanings they are implying by their use of certain terms. Guide them if you can to clarify, remove ambiguity, and become more aligned with industry-standard terms and meanings.

Configuration management and its partner process configuration control together keep a system and all of its elements managed in a cohesive, controlled way as changes, updates, or repair actions take place. Configuration management is a responsibility of both due care and due diligence and is vital to asset management. It is also a high-payoff set of process investments to make for improved information systems security. Configuration management ensures that the right stakeholders have made informed decisions to make changes, apply patches, or delete elements of your systems; configuration control ensures that those directed changes get made and that no other changes are allowed to take place.

Configuration management has perhaps the largest and most direct impact on an IT system's security posture. Without an active and effective configuration management and configuration control (CM/CC) system in place, your systems are essentially unmanaged and vulnerable. Consider as your starting point that straight from the shipping cartons, the default settings on all of your IT hardware, software, firmware, and data are often unsafe. One simple misconfiguration such as leaving a guest account open can bypass all other security controls. If by chance the new equipment or software you install is set up correctly and has no exploitable vulnerabilities still exposed, without configuration control, subsequent changes to that system and other systems it interacts or coexists with can re-expose those factory default weaknesses. Don't get the wrong impression here—without those factory or vendor default settings in place, you'd never be able to install and get the system up and running the first time. Once you do, of course, change them and lock them down tight.

The record-keeping that is the backbone of a good CM/CC system has another great payoff waiting for you, in the event that disaster strikes and you have to reload a bare-iron backup processing facility (or virgin VMs in the cloud) before you can get back into normal business operations. Those CM/CC records give you a known configuration baseline to use to *verify* that the backup images you loaded are configured, in all details, the way

your management processes said they should be—the way your live production systems had been configured just before disaster struck.

Your organization should start by developing a configuration management plan if it does not have one in operation already. A configuration management (CM) plan defines how an organization will manage the configuration of its hardware and software assets. It defines details such as the roles, responsibilities, policies, and procedures that are applicable. A configuration control board (CCB), which ITIL guidance refers to as a change advisory board (CAB), will manage the CM plan. As the CCB is comprised of qualified stakeholders from the organization, they will often be the authors, editors, reviewers, and approvers of the organization's configuration policies and procedures. They will also be tasked with applying and enforcing the CM plan and helping technical administrators adhere to and understand the CM plan. Most importantly, the CCB controls and approves changes throughout the lifecycle of the IT systems, which is why they may also be known as the change control board.

Configuration management and change control focus on the life history of individual configuration items and on sets of configuration items. A configuration item (CI) is one discrete part of an IT system, like a piece of hardware or software, that has configurable settings or parameters and should be under formal configuration control. A baseline configuration is a defined, desired set of configurations for a specific CI (or combine multiple CIs into an IT system), which has been formally reviewed and approved. A baseline configuration is valid for a given point in time and may need to be adjusted over time as software or hardware versions change, new vulnerabilities are discovered, or different usage and needs dictate the need for change. When the baseline configuration needs to change, it should be done only through predefined change control procedures. Deciding what a CI should be is a matter of perspective. Consider as an example that a modern platform system such as Microsoft Office Professional might contain between 5,000 to 10,000 individual files, or about 2 GB of code, configuration data, settings, forms, and templates. To the Microsoft Office developer team, each of those files is a CI. To your company's systems administrators who download licensed distribution kits, configure them, and install them onto dozens or hundreds (or more!) of endpoint systems throughout your company, they may see each new patch version of Office as one CI or see it as thousands of CIs (all those files and all of the patches to them). Fortunately, Microsoft (and many other platform product vendors) provide some pretty extensive maintenance management tools to help you manage their products as deployed systems, rather than as deployed swarms of a huge and unwieldy number of files.

Execute Change Management Process

As the systems security analyst and administrator, your duties may combine or overlap with those of other systems administrators who actually install, manage, and maintain

the operating systems, applications, platforms, web pages, and datasets that make up your organization's IT architecture. Without their extensive training and significant experience with those products, it's probably unrealistic for you to try to manage both the security configuration and the product configuration for each installed product. Let's look at a few of the methods and tools used in establishing and managing both kinds of configurations.

Manual configuration is the easiest to understand conceptually—it involves the administrator viewing and changing the configuration settings directly, either by editing a configuration settings data file or by using something like the Windows Registry Editor (regedit). Registry edits (or their equivalents in other operating systems environments) can also be done using batch or script files. Either way, this is a fine-grained, detailed, step-by-step process, which can be useful if you're stepping through various settings to diagnose a problem or as part of an incremental hardening process.

Configuration scanning tools can read the stored data structures used by the operating system and installed programs, extract information from those settings, and in some cases test some of those configuration settings for validity. The resulting list of all of these settings is sometimes called a *configuration enumeration*. NIST maintains a set of Common Configuration Enumerations that have been associated with security issues that are tracked in the National Vulnerability Database (NVD), and more recent versions of configuration scanning tools can help you detect similarities between a CCE and your system's current configuration. The CCE database can then provide you with insights and recommendations, drawn from best practices in the field, as to changes you should make in your systems to improve their overall security.

In the same breath, NIST and others often provide, specify, or recommend systems hardening information as it pertains to a given configuration enumeration. As a result, some professionals refer to the total bundle (the enumerated configuration and its related hardening information) as an *enumeration* or as a *set of hardening standards* for a particular configuration. Since the purpose of having the enumerated configurations in the first place is to collate hardening recommendations with specific configuration items and settings, this is to be expected. If in doubt as to what is meant or included, ask for clarification.

Another useful tool is a configuration change detection tool. It is different than a configuration scanner tool in that instead of asking the IT asset "Are you configured correctly?" it asks, "Did your configuration change?" It takes a snapshot of a given system's configurations, presumably after it was configured correctly and securely. Then, if any of the configurations are changed, it sends an alert to one or more relevant security stakeholders. Vendors are adding additional features and capabilities to both scanner tools and change detection tools, blurring the line between the two. Some tools now do both.

When you want to control how your security tools share data, you can use the Security Content Automation Protocol (SCAP). SCAP is a way for security tools to share data. It is an XML-based protocol that has many subcomponents called *specifications*, including one for CCE. It is a taxonomy for describing configuration requirements, which is essential because of the sheer number of configurations and their nuanced differences.

CCEs are written for, and are grouped by, specific IT products or technology types. The vulnerability equivalent to CCE is the Common Vulnerabilities and Exposures (CVE). CVE is more widely adopted than CCE because the vulnerability scanner market is larger and more mature than the configuration scanner market. In fact, some major vulnerability scanning tool vendors have added CCE (configuration) scanning to their traditional CVE (vulnerability) capabilities. Learn more about CCEs at `https://nvd.nist.gov/config/cce/index`.

In addition to other standards and guides, vendors (especially OS vendors) typically publish secure build outlines for their own products and often make tools available for provisioning and monitoring configurations.

Identify Security Impact

Any proposed change, even applying a patch kit or bug fix to alleviate a security problem, may inadvertently introduce a new vulnerability or a new risk into your systems and your business operations. Change packages should be examined to identify any potential changes to your operational procedures for getting work done with the affected systems and assets. Descriptions of the changes, and in particular the issues or vulnerabilities that are acknowledged as not addressed in the patch or update kit, should also be closely looked at to see if they suggest possible new areas of risks to your operations. If it's practical for you to delay installing the update until other organizations have installed it and operated on it for a short while, you may want to consider this—but only if you have an alternative way to protect your system from exploits targeted at the vulnerabilities the patch or update is going to remediate!

When analysis fails to surface anything to help alleviate your fears of causing more trouble and risk with an update than the fix is trying to eliminate, it may be time for some security-driven testing.

Testing/Implementing Patches, Fixes, and Updates

Chapter 7 goes into more detail on the overall software development process and the concepts behind the software development lifecycle (SDLC) models, both classic and cutting-edge, that are in widespread use today. As the security administrator or team member, you may need to be involved in the overall development process to ensure that any security-relevant issues, perspectives, functional requirements, and insights get incorporated into both the product as it is developed and the management process that keeps

that development on track. At some of those test opportunities—which there are more of in a large systems development than there would be for a small, tightly focused patch or update—security may need to be more of an active member of the test team and not just an interested stakeholder and observer. Your experience and insight about what happens when systems fail to be secure can be of great help to test teams as they conduct scenario-based test cases; your knowledge of how the application or system under test *should* be interacting with network and systems security monitoring and incident detection tools may also benefit the post-test analysis activities as well.

It is best and common practice to do security-related testing in an isolated testing environment, safely quarantined off from your live production environments. Virtual machines in tightly secured test and development subnets, and hosts are ideal for this. This contains any problems that the test may otherwise set loose into your production systems or out into the wild. It also allows you to be more aggressive in stressing the system under test than you could otherwise afford to be if testing were conducted on or associated with your live production environment.

You can also adapt penetration testing scenarios and approaches you would otherwise use against your systems hosted in an isolated testing environment, before you've released those new versions of the systems into live production and operational use. Black box, white box, or other forms of penetration testing may be quite useful, depending upon the nature of the changes you're trying to evaluate.

PARTICIPATE IN SECURITY AWARENESS AND TRAINING

In many respects, you, as the on-scene security professional, have the opportunity to influence one of the most critical choices facing your organization, and every organization. Are the people in that organization the strongest element in the defense, security, safety, and resiliency of their information systems, or are these same end users, builders, and maintainers of those systems the weakest link in that defense? This is not an issue of fact; it is a matter of choice. It is a matter of *opinion*. Shape that opinion.

Awareness is where you start shaping opinion, and in doing so, you inspire action—action to learn, action to become, action to change the way tasks get done and problems get set right. You might not be a trained and experienced educator, trainer, or developer of learning paths, course materials, and the tactics to engage your co-workers in making such an awareness campaign succeed. Don't worry about that. What you *can* and *should* do, as part of your professional due care and due diligence responsibilities, is engage with management and leadership at multiple levels to obtain their support and energy in moving in the right direction.

Increasing your co-workers' awareness of information security needs, issues, and opportunities is the first step. They'll then need a combination of the conceptual knowledge and the practical skills to translate that awareness into empowerment, and empowerment into action. Depending upon the lines of business your organization is involved in and the marketplaces or jurisdictions it operates in, there may be any number of risk management frameworks, information security policies and standards, or legal and regulatory requirements regarding effective security awareness, education, and training of your organization's workforce that must be complied with. This is not a cost or a burden; this is an opportunity for small, focused investments of effort to turn the tables on the threat actors and thereby take a significant bite out of the losses that might otherwise put your team out of work and the organization out of business.

Security Awareness Overview

It's easy to see that in almost every organization, no matter how large or small its workforce, no one single person can possess the knowledge, skills, abilities, and attitudes to successfully do all of the jobs that make that organization successful. By the same token, no one information security professional can keep all of the systems and elements of the IT architecture secure *and* plan, develop, and teach the security awareness, education, and training programs the rest of the workforce needs. What any of us *can* do—what *you* can do—is to take a thumbnail sketch of what such programs need to achieve, share this with management and leadership, and assist where you can with the expertise and talent you do have to make that sketch of a plan become reality. Let me offer you some thoughts about this, from my experiences as an educator, trainer, and information security professional.

Let's start with awareness—the informed recognition that a set of topics, ideas, and issues exists *and is important*. Awareness shines a different light on the day-to-day, triggering moments of recognition. Awareness shatters the false myths, the explanations that everybody "knows" but have never tested for validity. Simple but compelling examples can do this; even something as simple as "fake phishing" attack emails that you send to your own workforce can, over time, increase the percentage of that workforce that get better at spotting a possible attack and dealing with it immediately and correctly.

Education explains concepts and links them to awareness. Education can be formal, focused around an identified body of content or aimed at the student attaining a credential of some kind attesting to their accomplishment. Informal education can be just as effective and often is well suited to rapidly evolving situations. Education stimulates thinking and creativity. A short course in root cause analysis can start with getting students to recognize the power of simple, open-ended questions.

Training teaches skills and guides learners in becoming increasingly proficient in applying them to realistic situations. Training activities that use "spotters' guides," for example, can demonstrate packet sniffing and filtering or anti-phishing email screening techniques and then use checklist approaches as the frameworks of labs and exercises to enhance learners' abilities to recognize concepts in action and make informed decisions regarding actions to take.

Competency as the Criterion

It's well worth the investment of time and thought to create a short list of the key information security competencies that different subgroups of your workforce need, if they are going to be able to make real contributions to improving information security for the team as a whole. The larger your organization and the more diverse the individual workgroups are in terms of tasks, context, and the sensitivities of the information they work with, the greater the likelihood that you'll need numerous short lists of such competencies. This is okay; make this manageable by starting with the groups that seem to need even a small step-change in security effectiveness and work with them to identify these core competencies.

By the way, some education and training program professionals will refer to this core competencies approach as a *needs assessment*. The name does not matter; the results do. Both should produce as an outcome a list of tangible, clear statements of what learners need to learn and the standards by which they must be assessed to demonstrate the success of that learning.

It's likely that your company or organization has trainers and human resources developer talent within the HR or personnel department. Find them; get them involved. Get their help in translating these first few sets of core competencies into the next layer of detail: the activities that learners have to perform well at to demonstrate that they've successfully learned that competency to the required degree of rigor. Get them to help you find teaching and learning assets and materials that the company already has; or, get them to help you find other assets. Reuse what you can find, learning from how well it works, before spending the time to develop something custom-made for your situations, people, mission, and needs.

Build a Security Culture, One Awareness Step at a Time

You've successfully engaged others in the company to take on the tasks of selecting or developing the teaching and learning assets, structuring the courses, and finding the right people to act as trainers and teachers. You've got them managing the identification of which employees need what levels of learning, how often they need it, and when they need to get the learning accomplished. As the on-shift or day staff security administrator, that's a great segregation of duties to achieve! Now what?

Walk the hallways of the company's campus or locations; keep your eyes and ears open for signs that awareness, learning, and skills-building are happening. Look for signs of trouble that suggest it isn't working fast enough or well enough. Step into those situations informally and casually, and lead by example and inspire by action and word. Suggest to people in these problematic contexts, be they workers, supervisors, or mid-level managers, that they've got the opportunity to empower themselves, and you can help them.

Too many organizations fall into the administratively simple task of regularly scheduling repetitive training activities. These could be messaging opportunities that strengthen each worker's future with the company by enhancing the organization's survival and success. Instead, they oftentimes turn them into tick-the-box, square-filling exercises in futility. If this is happening in your organization, shine some light on it; help others become aware of the need to turn that messaging around. Quickly.

PARTICIPATE IN PHYSICAL SECURITY OPERATIONS

Information security specialists, such as SSCPs, need to be aware of all threats to the information systems in their care and be able to assist, advise, and take action as required across many functional areas in their organization. If your company is truly cloud-based, with no data center of its own, you've still got threats in the physical domain to contend with. Remember, too, that your attacker could turn out to be an insider who turns against your team for any number of political, financial, emotional, or personal reasons.

Physical Access Control

If the attackers can get to your systems, they've got a chance to be able to get *into* them. This starts in the physical domain, where access includes physical contact at Layer 1 network systems, at the USB ports or memory card slots on your endpoints and other devices. It includes being able to see the blinking LEDs on routers (which blink with each 1 or 0 being sent down the wire), and it includes being bold as brass and just walking into your office spaces as if they're a pizza delivery person or business visitor. And although we've not yet seen it reported, it won't be long now before we do see an attacker using hobbyist-grade UAVs to carry out intrusion attempts.

Chapter 2 will look at the concept of defense in depth, integrating a variety of deterrence, prevention, and detection capabilities to defend the points of entry into your systems. Threat modeling, done during the risk assessment and vulnerability assessment phases (which Chapter 3 examines in more detail), have given you maps of your systems architecture, which show it at the data, control, and management planes as well as in the physical dimension. Start at the outermost perimeter in those four planes and put on your penetration-tester hat to see these control concepts in action.

One major caution: What you are about to do is tantamount to penetration testing, and to keep that testing ethical, you need to first make sure that you're on the right side of law and ethics. Before you take *any* action that might be construed as an attempted penetration of an organization's information systems or properties under their control, gain their owners and senior managers permission *in writing*. Lay out a detailed plan of what you are going to attempt to do, why you propose it as worthwhile, and what you anticipate the disruptions to normal business operations might be. Work with them to specify how you'll monitor and control the penetration test activities and how you'll suspend or terminate them immediately if required. As you learn with each step, err on the side of caution and go back to that management team and ask for their written permission to take the *next* step.

At a minimum, this will keep you out of jail. It will enhance your chances of staying employed. It will also go a long way toward increasing the awareness of the threat and the opportunity that your management and leadership stakeholders have to *do something* about it.

✔ Don't Fail to Imagine

The vast majority of businesses and nonprofit organizations have almost nothing to do with national defense or with international intrigue; their leaders, managers, and owners see themselves as light years away from international terrorist plots or organized crime. And they probably are. Unfortunately, this distance can bring a false sense of security with it, one that turns off their imagination.

In virtually every cyber attack, the target is the *data* that the organization holds. Data about their employees, their customers, or their suppliers; or transaction histories with their partners and their banks. Attackers may have far more reasons for finding value in *your* data than you think.

Without your data, you can't operate. With your data, your attackers can gain in ways you don't have to imagine in order to stop cybercrime in its tracks.

Property Approach

From early reconnaissance and target selection onward, an APT actor will need to see, sense, observe, and probe at your facilities, your people, and your IT systems. You need to balance allowing these contacts for *legitimate* outsiders while not making it too easy for

a hostile agent to learn too much. You don't control the Internet any more than you control the physical spaces outside of the property line around the buildings your company occupies, but you can and should consider what you choose to make visible, audible, or otherwise physically observable, for example, via:

- Visual line of sight, depending on the sensitivity of the organization's operations. Line of sight might be obscured by limiting windows in construction, covering windows in sensitive areas, obstructing views with landscaping/formation, or other means.
- Vehicular approach, including roads and driveways toward the property/facilities. For secure facilities, these should deter a straight approach to disallow a drive to build up excessive speed and should include obstacles with bollards, barriers, or retractable tire spikes.
- Movement patterns of your workforce can reveal when they're working a special, important activity that demands a surge of effort, versus a normal routine pattern of arrivals and departures.

In the digital domain, use periodic black-box ethical penetration testing techniques to examine all publicly-facing information that your organization makes available on web pages, via e-commerce or e-business connections, and even in advertising and print media. Port scanning and network mapping also may show you spots where your systems reveal too much about themselves.

Perimeter

At the outer boundary of the property, security controls can be implemented for access control.

- **Fences/walls:** While generally seen as deterrent or preventive controls, fences and walls can also be combined with additional mechanisms to offer detection capabilities.
- **Cameras:** Cameras serve a deterrent purpose but can be combined with monitoring capabilities (such as guards watching a video feed or motion sensors) for detection functions. Know that it's fairly easy for dedicated attackers to separate the cameras that are actually monitored from those that are "perimeter dressing" and most often ignored.
- **Buried lines:** While these serve no deterrent function, underground sensors can be used for intrusion detection within the border of a property.
- **Access control points:** Guard stations or gates can be staffed or equipped with additional mechanisms (card readers, cameras, turnstiles, etc.).

- **Patrols:** Guards (human or canine) can provide deterrent, detective, corrective, and recovery controls.
- **Motion sensors:** There are a variety of technologies that support the organization's ability to surveil the perimeter and any area outside facilities, including the cameras and buried lines, as well as microwave, laser, acoustic, and infrared systems.
- **Lighting:** Well-lit areas serve both deterrent and detective purposes. Continual maintenance of all lighting sources is crucial, as a burned-out or broken bulb can defeat any security benefit the light might provide.

Parking

The most dangerous workplace location is the site where vehicles and pedestrians meet. It is imperative to include sufficient lighting, signage, and conditions (width of right-of-way, crosswalks, etc.) to minimize the possibility of threats to human health and safety. Monitoring is also useful, as parking areas are often locations that are accessible to the public and have been frequently used to stage criminal activity (workplace violence, robbery, rape, murder, etc.).

If the parking structure allows for entry to the facility, this entry should be equipped with access controls, and all entryways should feed to a single reception point within the facility.

Generators and fuel storage, as well as utility access (power lines, water/sewer pipes, etc.), should be protected from vehicular traffic, either with distance or with additional physical obstructions. There must be sufficient access for fuel delivery traffic, but this should be severely limited to reduce risk.

Facility Entrance

In addition to the other entrance controls already mentioned, the entry to the facility might include the following:

- **Reception staff:** This includes guards or administrative personnel who observe people entering and leaving the facility.
- **Logging:** This may be as technologically rudimentary as a sign-in book or combined with sophisticated badging/monitoring capabilities.
- **Flow control:** Turnstiles or other mechanisms ensure only one person at a time can pass, typically only after presenting a credential (such as a badge or biometric element).

Internal Access Controls

In addition to the other access control elements used for maintaining physical control of the workplace environment listed elsewhere in the book, the security practitioner should be familiar with the following:

- **Safes:** Secure containers that can offer protection from unauthorized access, fire, water damage, and, in some cases, chemical contaminants. Both the safe itself and the lock on the safe should be rated by a standards body for specific criteria, according to the particular needs of the organization.
- **Secure processing areas:** Specific areas within the workplace that are set aside, both administratively, technically, and physically, from the rest of the production environment. These are typified by secure entryways, severe limitations on personnel access, hardened structures (walls, no windows, etc.), and electromagnetic shielding. In the U.S. government sphere, these are referred to as *sensitive compartmented information facilities* (SCIFs), although the term has begun to see wider use in nongovernment activities in recent years.

TIP Can Visitors Spot your Vulnerabilities? "Reconnaissance by walking around" is a time-honored component of many an intrusion; it's even easier nowadays when smartphones can conduct full Wi-Fi surveys. Try it yourself, as part of an ethical penetration test.

The Data Center

As the focal point of the data assets of the organization, the data center is in particular need of protection within the property/facility. The data center also has some specific requirements that make it somewhat different than the rest of the production environment. In addition to the other access controls placed on secure areas within the workplace (discussed earlier in this chapter and in Chapter 5), security of the data center should include consideration of the following factors:

- **Ambient temperature:** IT components generally function better in relatively cold conditions; if the area is too hot, the machines will not function optimally. However, if the area is too cold, it will cause discomfort for personnel.
- **Humidity:** An interior atmosphere that is too dry will increase the potential for electrostatic discharge. An atmosphere that is too damp will increase the potential for development of mold, mildew, and insects.

Standards for maintaining a desirable range of data center environmental conditions should be used to establish targets. One such reference is the ASHRAE Technical Committee 9.9 thermal guidelines for data centers; see `http://ecoinfo.cnrs.fr/IMG/pdf/ashrae_2011_thermal_guidelines_data_center.pdf`.

The data center should also be designed, constructed, and equipped for resiliency, such that it is resistant to unplanned outages from human error/attack, system/component failure, or natural effects. This is typically accomplished by including a great deal of redundancy within the data center. The use of design standards to achieve a significant level of robustness and resiliency is highly recommended.

The Uptime Institute publishes a multitier standard for use by data center owners in determining and demonstrating their particular requirements and capabilities ("Data Center Site Infrastructure Tier Standard: Topology"; see `https://uptimeinstitute.com/tiers`). The tiers range in purpose and requirements from basic data centers that might be used for archiving or occasional data storage to facilities that support life-critical processes. The CISSP should have a cursory knowledge of the four-tier levels and their descriptions. (For more information, see `https://journal.uptimeinstitute.com/explaining-uptime-institutes-tier-classification-system/.`)

The standard is free for review/guidance; certification against the standard is performed only by the Uptime Institute and requires payment.

Organizations that receive Uptime Institute tier certification for their data centers can be listed in the Institute's online register: `https://uptimeinstitute.com/TierCertification/allCertifications.php?page=1&ipp=All`.

Finally, fire poses a significant, common risk to data centers because of the high potential for occurrence and because of the disproportionately heavy impact a data center fire would have on the organization. The selection, design, implementation, maintenance, and use of fire protection and alarm systems can be quite complex, and in many jurisdictions must be undertaken by a properly licensed fire protection engineer. Municipal standards such as building codes also must be taken into account. Insurance providers may also levy strict inspection and compliance constraints on any and all fire protection systems and practices in order to maintain policy coverage. This all goes well beyond what the SSCP can or should attempt to take on.

Service Level Agreements

In the modern IT environment, there are many reasons (not the least of which is cost) for an organization to consider contracting with an external service provider to handle regular operational tasks and functions. To create a contract favorable for both parties in this sort of managed services arrangement, everyone involved must clearly understand what is being requested, what is being provided, what the cost is, and who is responsible for what. This is particularly important in what could be considered the most popular current form of managed services: cloud-managed services. In the majority of cloud-managed service contracts, the cloud provider and customer must determine the expected level of service, and the contract or service level agreement is the element that gives both parties the confidence to expect defined outcomes: assuring the provider that they will receive payment and assuring the customer that the service will meet the customer's needs.

In these cases, you need a formal agreement that defines the roles and responsibility of each party, explicit to the point where it can be easily understood and measured. The common name for this is the service level agreement. However, depending on the services provided, the agreement can go by other names, like *network services agreement*, *interconnection security agreement*, etc. The SLA is part of the overall contract but deals directly with the quantifiable, discrete elements of service delivery.

These are scenarios where an organization might need an SLA:

- Third-party security services
  - Monitoring/scanning
  - Security operations center/response-type services
  - Media courier/media disposal
  - Physical security
- Hosted/cloud
  - Servers
  - Storage
  - Services
- Interconnecting information systems, especially with data feed/pull/push
- Supply chain scenarios

The SLA portion of the contract vehicle is best limited to those elements of the managed service that are routinely provided as part of continual operational requirements; the SLA is not the optimum place for including contingency requirements (such as BCDR tasks) or for anything that cannot be distilled into a numeric value.

Specific Terms and Metrics

To be effective (and enforceable), an SLA must use clear and unambiguous language to specify its terms and conditions for all services that each party brings to the contract. Key performance indicators or other quality of service metrics should also be defined in the SLA, along with explanations as to how they are measured, computed, and reported. Without this, there is no basis for measuring or knowing whether a provider is providing the agreed level of service.

Amazon Web Services (AWS), a well-known cloud service provider, uses a standard SLA for their Elastic Cloud Compute (EC2) services, which you can review at `https://aws.amazon.com/ec2/sla/`. Among other items, it specifies a server uptime metric:

- If your servers enjoy anything above 99.99 percent uptime, AWS has met its SLA.

- If your servers have anywhere between 99.00 and 99.99 percent uptime for the month, you will get a 10 percent discount on the service fee for that period.
- For anything less than 99.00 percent, you will get a 30 percent discount for your hosting for that month.

This is a good example not only because the metrics and terms are clear but also because it is clear about what happens in the event of noncompliance with the SLA. The contracting manager (in conjunction with the organization's IT department) must determine whether the price reduction would realistically offset the loss in productivity a service outage would cause; if the cost of the outage outweighs the benefit of the rebate/discount, the SLA is insufficient for the customer's needs.

Mechanism for Monitoring Service

It is not enough, however, just to understand the terms of the SLA. You also need a mechanism with which to monitor and measure whether the service provided matches the level specified in the SLA.

To continue with the previous example of AWS, visit `https://status.aws.amazon.com/`. You will initially see a dashboard similar to Figure 1.3. The horizontal rows represent the AWS regions. If you look at the corresponding region where your servers are hosted, you can see whether they are having, or have had, any degradation of service or outages.

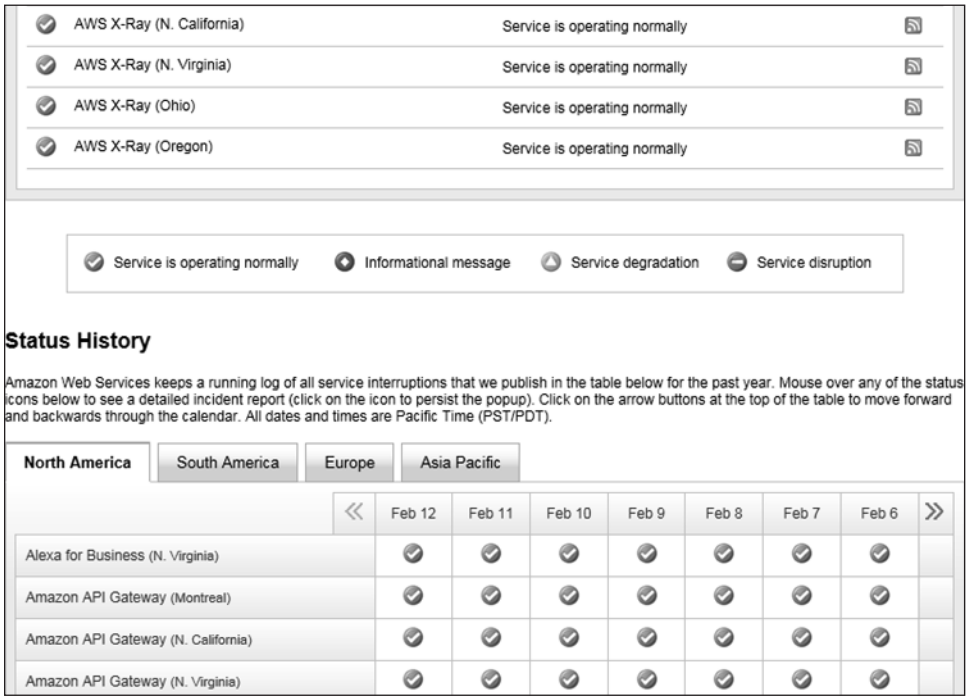

FIGURE 1.3 **AWS dashboard**

While this dashboard can be used to inform the customer as to the efficacy of the service overall, it might not provide, by itself, the level of assurance the customer desires; the information is necessarily coming from the provider, and the provider has a vested interest in the outcomes of the data (i.e., getting paid) and so is inherently biased. For such SLA elements, the customer may prefer some third-party validation of the service/data to feel confident that the reporting mechanism adequately reflects the actual level of service provided/received.

SUMMARY

It's in the day-to-day details that you have the greatest opportunity to thwart an attacker from gaining meaningful insights about your information systems and then leveraging those insights to attempt an intrusion. It's in the day to day that you mentally, virtually, and physically patrol your perimeters, layer by layer, and stay in touch with the sensors and preventers that are working to keep things safe and secure. It's hard to keep a paranoid edge to your awareness; it's hard to avoid being lulled into a no-news-is-good-news complacency. One built-in advantage you have is that in a properly planned and executed security posture, the list of things you need to check up on is almost limitless: Boredom should never be your problem! Get curious, and stay curious, as you check with the badge readers and the other AAA elements of your access control technologies. Review what the security information logging and analysis systems are trying to tell you. Touch base with the help-desk people, with visitor control, and with all of the human elements that make your security strong—or break it, if left ignored and uncared for.

Making information security into something that works effectively every day, every hour, is an operational and administrative task. It needs you to manage it by physically and virtually walking around. Think like a hacker; turn that thinking into ideas for ethical penetration testing, even if only on paper or sitting around a conference table with people from other functional areas in your organization. Hear what they say, and help them grow the security culture you all need to enjoy and be safe in.