SECURITY MANAGEMENT

A **CRCPRESS** FREEBOOK

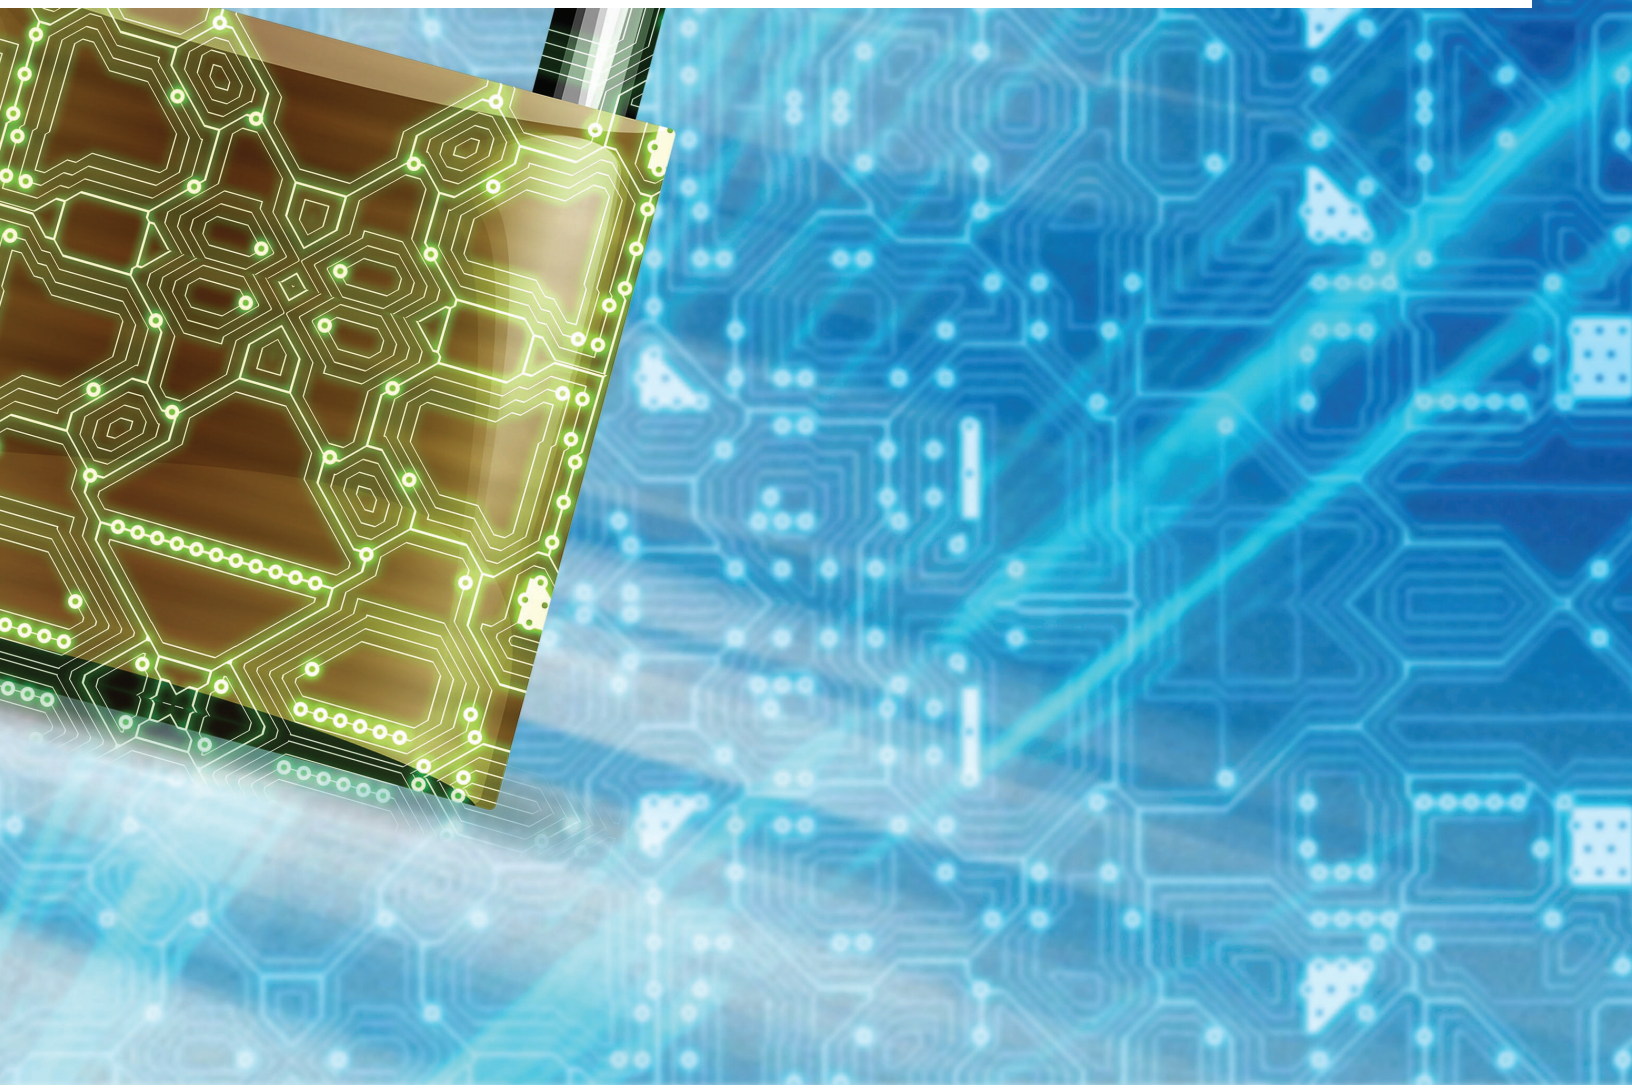

CRC Press
Taylor & Francis Group

TABLE OF CONTENTS

READ THE LATEST ON SECURITY WITH THESE KEY TITLES

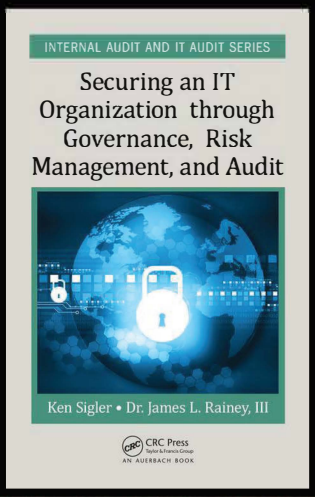

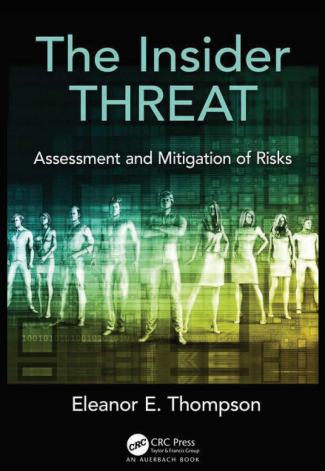

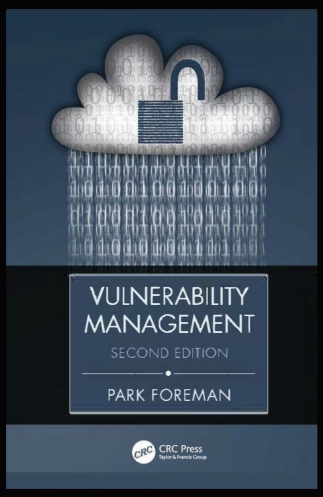

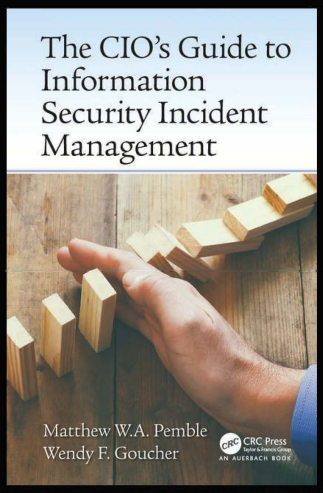

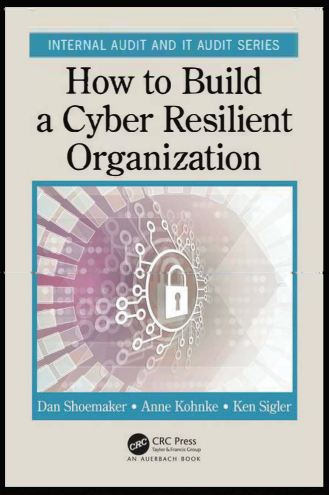

Introduction

This FreeBook contains introductions to topics essential to Security Management such as Cybersecurity and Crisis Management.

Securing an IT Organization through Governance, Risk Management, and Audit introduces two internationally recognized bodies of knowledge: Control Objectives for Information and Related Technology (COBIT 5) from a cybersecurity perspective and the NIST Framework for Improving Critical Infrastructure Cybersecurity (CSF).

The Insider Threat book provides emergent knowledge relating to physical, cyber, and human risk mitigation in a practical and readable approach for the corporate environment. It presents and discusses practical applications of risk management techniques along with useable practical policy change options.

Vulnerability Management demonstrates a proactive approach to the discipline. Illustrated with examples drawn from Park Foreman's more than three decades of multinational experience, the book demonstrates how much easier it is to manage potential weaknesses than to clean up after a violation.

The CIO's Guide to Information Security Incident Management will help IT and business operations managers who have been tasked with addressing security issues. It provides a solid understanding of security incident response and detailed guidance in the setting up and running of specialist incident management teams.

How to Build a Cyber-Resilient Organization presents a standard methodology approach to cyber-resilience. Readers will learn how to design a cyber-resilient architecture for a given organization as well as how to maintain a state of cyber-resilience in its day-to-day operation.

CHAPTER

1

CYBERSECURITY RISK MANAGEMENT

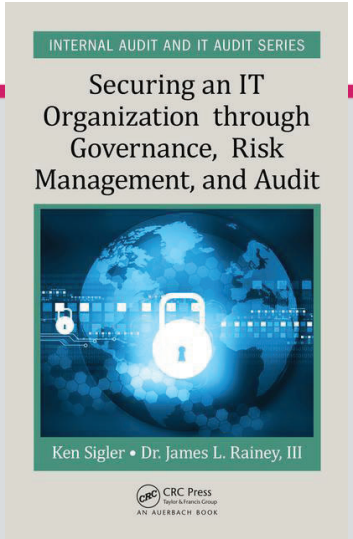

This chapter is excerpted from

Securing an IT Organization through Governance, Risk Management, and Audit

by Ken E. Sigler, James L. Rainey, III

Learn more

1

Cybersecurity Risk Management

After reading this chapter and completing the case project, you will

- Understand the definition of *cybersecurity* and how it fits into the overall scope of information and communications technology (ICT) organization management;
- Understand the importance of cybersecurity risk management as it relates to managing ICT life cycle processes;
- Understand the relationships between managing ICT governance, risk, and audit;
- Understand the best practices for using standards and frameworks to foster productivity in managing ICT governance, risk, and audit; and
- Understand the benefits of using standards and frameworks within ICT life cycle processes.

Cybersecurity

This book is based on a simple reality. The failure to include governance, risk management, and audit within the management structure of an information and communications technology (ICT) operation leads to unreliable and insecure products. Accordingly, this book describes an approach that lets ICT managers establish and sustain a logical and secure management process. The advice shared in the book is supported by a well-defined and standard set of management practices that have been proved to ensure trustworthy products.

A common cliché used by many ICT managers is, "if it isn't broke, why fix it?" The answer lies in the growing number of harmful effects of exploitation. As ICT systems continue to take advantage of the Internet and cloud computing technologies, those

systems continue to grow exponentially, connecting layer of software made up of trillions of lines of code. Those layers impact every aspect of the general public's way of life, from social networking and other forms of personal entertainment to national defense. A security attack in any of those layers could lead to personal tragedy or national disaster. How serious is the problem? Veracode, a major ICT security firm found that "58 percent of all software applications across supplier types [failed] to meet acceptable levels of security" (Veracode 2012).

Cybersecurity: A Definition

The field of cybersecurity has taken on a number of definitions over the years. However, the common body of knowledge throughout the ICT industry agrees that cybersecurity is concerned with creating and implementing processes that identify emerging threats in addition to providing cost-effective countermeasures to address those threats. The National Initiative for Cybersecurity Careers and Studies more formally defines *cybersecurity* as "strategy, policy, and standards regarding the security of and operations in cyberspace, and encompass[ing] the full range of threat reduction, vulnerability reduction, deterrence, international engagement, incident response, resiliency, and recovery policies and activities, including computer network operations, information assurance, law enforcement, diplomacy, military, and intelligence missions as they relate to the security and stability of the global information and communications infrastructure" (National Initiative for Cybersecurity Careers and Studies 2014). Cybersecurity as a discipline has grown quickly, most notably since the 9/11 attacks. Since that time, given the critical role that the Internet plays in our lives, a formalized discipline to study effective ways to ensure confidentiality, integrity, availability, authentication, and nonrepudiation of digital information has moved to the front of our national priority list.

Despite the national attention the field of cybersecurity has received since the beginning of the 21st century, there is still much debate about what mechanism provides the right set of actions to eliminate, or at least minimize, security attacks. Part of the reason for this debate is

that the field of cybersecurity consists of contributions from a number of different disciplines. These disciplines include the following:

- Traditional computer security and computer science studies, which provide the knowledge for safeguarding electronic information
- Networking studies, which supply the knowledge necessary to secure storage and transmission of data and information
- Software engineering, which adds process considerations, such as verification, validation, configuration management, and other forms of life cycle process security
- Business management, which provides the knowledge necessary to conduct project management, create and sustain security policy, and enforce contract and regulation compliance
- Legal studies, which contribute consideration of intellectual property, privacy rights, copyright protection, cyberlaw, and cyberlitigation

All of these disciplines provide a distinct contribution to the underlying effort of protecting an organization's data and information infrastructures. It is only through collaboration between all of these areas, working together toward goal, that significant organizational effectiveness can be achieved toward establishing best practices for cybersecurity. There is still an issue to resolve, however. Many ICT organizations still struggle with managing "who does what." There remains a lack of clear understanding regarding the scope of contribution each discipline serves within the cybersecurity initiative. Regardless of the management structure, common sense tells us that there are three vital components that need to exist within an organization in order for cybersecurity to be achieved: ICT governance, controls, and audit. In a later section of this chapter you will learn how implementing a framework based on those practices is vital to the success of cybersecurity initiatives in an organization.

INSIGHT Cybersecurity Myths That Small Companies Still Believe

High-profile breaches at Target (TGT), Home Depot (HD), and JPMorgan Chase (JPM) have put cybersecurity on the agenda for

companies large and small. But despite the ongoing media commentary and "best practices" memos, consultant Adam Epstein of Third Creek Advisors notes that board members of small-cap companies and those considering or preparing initial public offerings are still befuddled by persistent myths on this topic.

The confused companies include many in Silicon Valley, where one would expect to find more tech savvy, he says. I asked Epstein, the author of a how-to book for corporate boards, to bang out a primer on what directors think they know about cyber threats but really don't. Herewith, his free advice:

1. **Cyber breaches are preventable.** No, they're not. Breaches are a matter of when, not if. As security guru Tom Ridge recently noted in my interview with him in *Directorship* magazine, your networks have likely already been breached. If Fortune 50 companies with nine-digit annual cybersecurity budgets can't prevent breaches, neither can you. Effective cybersecurity is more about identifying corporate "crown jewels," making it as difficult as possible for them to leave the building, and having a thoughtful plan for post-breach resilience.
2. **The information technology (IT) team is on it.** No, probably not. Boardroom cybersecurity oversight generally consists of inviting the head of IT to make a periodic presentation on the company's firewalls and antivirus software. Lacking security experts, most boards collectively exhale on hearing the IT update. Unfortunately, cybersecurity is only partially an IT issue. It's also a matter of corporate culture, employee training, and physical security. You need to worry about disgruntled employees and your supply chain, not to mention that little company you just acquired. That's way beyond IT.
3. **Cyber theft is about credit cards.** In the past several months, I've consulted with several boards whose members said that because their businesses don't store or process credit card data, this area isn't a cause for concern. Wrong. Cyber thieves have disparate goals, ranging from semi-benign mayhem, to espionage, to misappropriation, to terrorism. Credit card information is certainly a target, but so is personal info, intellectual

property, strategy memos, customer lists, and other nonpublic information.

4. **Always disclose cyber incursions immediately.** While it's admirable to want to get out in front of breach incidents and voluntarily disclose them, this can sometimes put a board at a disadvantage. Consider the Target breach, where the size and nature of the crisis expanded substantively with each press release. Malware can morph after being detected and wreak further havoc. It's often unlikely that the first information received by the board about a breach will be accurate and comprehensive, so exercise caution not to complicate a crisis by voluntarily misrepresenting it.
5. **No worries, we've got insurance for this.** A lot of so-called cyber coverage results from a three-page application that barely addresses the quality and extent of your company's computer-network architecture, physical and data security protocols, and corporate risk culture. The resulting coverage usually comes up short. Scores of cyber policies exclude more than they cover. Make sure the policy is underwritten after extensive, informed security assessments of your company—not just a standardized form sent via e-mail.

Good luck. You'll need that, too. (Barrett 2014)

Cybersecurity Risk Management

Simply put, risk management is the practice of looking at what could go wrong and then deciding on ways to prevent or minimize potential problems. It encompasses four components: frame, assess respond, and monitor. We all carry out informal risk management numerous times in the course of a day without even realizing it. Every time we cross a street, we stop to weigh the risk of rushing in front of oncoming traffic, waiting for the light to change, using the crosswalk, etc. Our ability to analyze the consequences of each decision is risk assessment. What we decide to do after performing that quick analysis is risk response based on proper early training and our experience of crossing a road. We may decide to wait for the traffic light and use the crosswalk, which greatly reduces the potential risk; we may follow

someone else across the street, allowing them to make the decision for us; or we may simply choose not to cross the street. These decisions are a result of our risk assessment of the situation. If you make it across the street, you remember what worked. If anything went wrong, such as a honked horn or brakes squealing, you should evaluate (or monitor) whether another choice would have been better.

For a manager the issues of risk assessment may seem difficult and the right decisions for risk management may seem challenging; but the principles remain the same. It is the responsibility of management to make the best decision based on the information at hand. A well-structured risk management process, when used effectively, can help. In the case of a local government, for example, a citizen may report a pothole on a local road and you are obligated to determine an appropriate response. There are many factors to consider: what if a car gets damaged driving on the pothole? Is the cost to fix the pothole justified by the potential consequences? What if a citizen sues or seeks restitutions for the damage caused by the pothole? You have to analyze the risk and then decide how to manage the problem. Is it best to put signs around the pothole warning citizens? Should you pay overtime to send a road crew out to fix it? Do you ignore the problem? Risk assessment allows managers to evaluate what needs to be protected relative to operational needs and financial resources. This is an ongoing process of evaluating threats and vulnerabilities and then establishing an appropriate risk management process to mitigate potential monetary losses and harm to an organization's reputation. For cybersecurity, the program should be appropriate for the degree of risk associated with the organization's systems, networks, and information assets. For example, organizations accepting online payments are exposed to more risks compared to websites with only static information.

Risk Management Components

Risk management and a risk management framework seem to be the same thing, but it is important to understand the distinction between the two. The risk management process is specifically detailed by the National Institute of Standards and Technology (NIST) in three different volumes. NIST SP 800-30, "Guide for Conducting Risk Assessments," provides an overview of how risk management fits into

the system development life cycle and describes how to conduct risk assessments in addition to how to mitigate risks. NIST SP 800-37 discusses the risk management framework defined by NIST. Finally, NIST SP 800-39, "Managing Information Security Risk," defines the multitiered, organization-wide approach to risk management that is discussed in this chapter.

Managing risks is a difficult, multidimensional activity that requires contributions from everyone within the ICT organization. The management at the top tier has the responsibility of providing the strategic vision and ensuring that the goals and objectives for the organization are met. The middle management plans, executes, and manages projects. Individuals that middle managers oversee have roles that require them to operate the information systems supporting the organization's missions/business functions. NIST characterizes risk management as "a comprehensive process that requires organizations to: (i) frame risk (i.e., establish the context for risk-based decisions); (ii) assess risk; (iii) respond to risk once determined; and (iv) monitor risk on an ongoing basis using effective organizational communications and a feedback loop for continuous improvement in the risk-related activities of organizations. Risk management is carried out as a holistic, organization-wide activity that addresses risk from the strategic level to the tactical level, ensuring that risk-based decision making is integrated into every aspect of the organization" (Ross 2011). Figure 1.1 shows the correlation between the components of

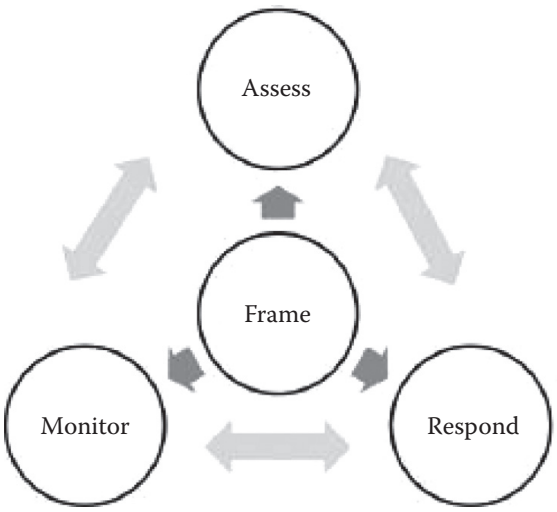

Figure 1.1 Risk management components.

risk management. Each of them interrelated with one another and lines of communication go between them. The output from one component becomes the input to another component. It should be noted that not every organization or periodical will use the same terminology for the four components. However, the criteria from which the components are based, and are described in the rest of this section, tend to remain consistent regardless of the terminology used for each component.

The first component of risk management addresses how ICT organizations frame risk. The senior management within the organization establishes the risk framework that will be used to define risk assumptions, risk constraints, risk tolerances, and risk priorities. Defining risk assumptions includes determining the likelihood that a vulnerability, threat, or occurrence could impact the organization and what the consequences or impact would be if it were to occur. Issues in the enterprise that restrict or slow risk assessments, risk response, or risk monitoring are categorized as risk constraints. Risk tolerances are those possible events or occurrences whose impacts on the organization are acceptable; often these risks are deemed acceptable because of the excessive cost of countering them. Finally, risk priorities are those events that must be protected against and systems that have a reduced risk tolerance. Many organizations prioritize system risk acceptance based on whether or not the systems support critical business or mission functions, as these systems have the lowest risk tolerance and highest risk priority.

The second component puts into context the organization's practices to assess risk based on the organizational risk frame. Before the organization commits resources to cybersecurity and ICT controls, it must know which assets have protection and the extent to which those assets are vulnerable. Risk assessment helps to answer those questions and determine the most cost-effective set of controls for protecting assets. Important to note is that not all risks can be anticipated and measured, but most organizations are able to gain understanding of the risks they face. Through risk assessment, managers try to determine the value of information assets, points of vulnerability, the likely frequency of the problem, and the potential for damage. For example, if some form of cybersecurity event is likely to occur no more than once a year, with a maximum of a $1000 loss to the organization, it

might not be justifiable to spend $20,000 on the design and maintenance of a control to protect against that event. However, that same event could be found to occur once a day, with a potential loss of $300,000 a year. In that case, $100,000 spent on a control might be appropriate. Once the risks have been assessed, ICT management can concentrate on the organizational hardware and software assets (or control points) with the greatest vulnerability and potential for loss.

The third component provides practices related to how organizations respond to risk once it is identified. This identification normally is an input to the risk response from the risk assessment component in the form of the determination of risk, but it can also come from the risk frame in the form of the risk management strategy. The risk response serves to provide an organization-wide, consistent response that addresses the risk frame. This includes developing courses of action, evaluating alternative courses of action, determining the appropriate course or courses of action, and implementing the risk response based on the selection. These steps are illustrated in Figure 1.2. The selection made has the potential to change the organization's risk procedure and, once made, the other components of the risk management process need to be evaluated for necessary changes.

The fourth and final component of risk management is concerned with how organizations monitor risk over time. This component validates that the risk program has implemented the planned risk response and that information security plans are derived from traceable mission/business functions. It also determines the effectiveness of ongoing risk response plans and determines and identifies changes

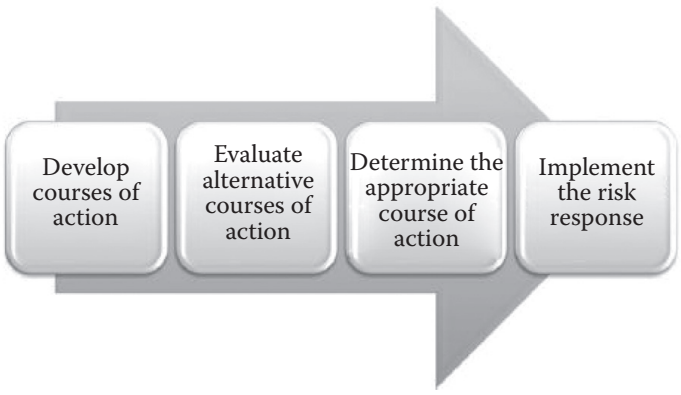

Figure 1.2 Steps of risk response.

in the environment that will impact the risk profile of the organization. The risk program can be modified as needed to respond to changes identified in the monitoring process. These changes initiate updates to the organization's risk assessment, risk response, and risk frame components.

As indicated in the four components of risk management that have been described, organizations must also consider external risk relationships when necessary. These external entities include those in which there is an actual or potential risk relationship. For example, consideration needs to be given to organizations that could impose risks on, transfer risks to, or communicate risks to other organizations, as well as those to which organizations could impose, transfer, or communicate risks. Depending on the type of business being assessed, external risk relationships could include suppliers, the customer's business partners, and service providers. For those organizations that have already identified persistent threats, specific attention should be given to the risk posed by suppliers within the organizations supply chain. Although management has control over the risks only within the boundaries of their organization, the more that an organization is aware of external risks, the easier it will be to implement internal practices to safeguard the potential for unexpected event caused by information sharing.

Risk Management Tiered Approach

In order to implement the risk management process throughout the organization, a three-tiered approach can be is used to manage risk at the

1. Organization level;
2. Mission/business process level; and
3. Information system level.

The risk management process is carried out through the three tiers with an underlying objective of continuous risk-related process improvement throughout the organization and effective communication between the tiers and among all stakeholders with a shared interest in the mission/business success of the organization. Figure 1.3 shows the three-tiered approach to risk management along with some of its key characteristics.

Figure 1.3 Three-tiered risk management approach.

Tier 1: Organizational Level

A risk management program is not going to be successful unless strategies are implemented and properly managed. Tier 1 is focused on risk from an organizational perspective by establishing and implementing governance structures that are consistent with the strategic goals and objectives of organizations and the requirements defined by federal laws, directives, policies, regulations, standards, and missions/business functions. The Control Objectives for Information and Related Technology (COBIT 5), a leading framework for ICT governance and management, and introduced in detail in Chapter 8, defines governance as follows: "Governance ensures that stakeholder needs, conditions and options are evaluated to determine balanced, agreed-on enterprise objectives to be achieved; setting direction through prioritization and decision making; and monitoring performance and compliance against agreed-on direction and objectives" (Information Systems Audit and Control Association 2012). Essentially, tier 1 implements the first component of risk management, risk framing, by providing the governance for all risk management activities performed through the organization. Tier 1 risk management activities directly affect and serve as a basis for the success of the activities performed at tier 2 and tier 3. For example, tier 1 defines the missions

and business functions of the organization that, in turn, influence the design and development of the mission/business processes created at tier 2 to accomplish those missions/business functions. Tier 1 also provides a prioritization of missions/business functions, which, in turn, drives investment strategies and funding decisions, which, in turn, affect the development of ICT system (including cloud computing infrastructures and embedded cybersecurity architecture) at tier 2 and the allocations and deployment of management, operational, and technical security controls at tier 3. Trails of feedback such as these could, and often does, result in changes to the organizations risk framework.

Tier 2: Mission/Business Process Level

Tier 2 addresses risk from a mission/business process perspective. The input to this level is the risk context, decisions, and output of the activities at tier 1. This level in the organization has specific goals to ensure that the organization remains viable. Human resources is a good example of a mission or business process level function of an organization. Activities typical at tier 2 include the following:

- Defining the mission of business need
- Prioritizing the mission or business processes
- Defining the type of information required to carry out the mission or business processes
- Incorporating and establishing ICT solutions with required security components

Enterprise architecture decisions made at tier 2 determine the acceptable technology solutions that can be implemented at tier 3. Additionally, decisions made by management at tier 2 influence the allocation of certain security controls to specific components or information systems once they are implemented at tier 3, based on the organization's established information needs. Managers at this level may determine what technologies are acceptable for processing the information created from a specific business function. To achieve success, the selection of approved and prohibited technologies should be well documented and distributed to the business function area and all stakeholders (including ICT developers and administrators) who support the business function.

Tier 3: Information System Level

Tier 3 contains the activities associated with risk from an information system perspective and is guided by the risk context, risk decisions, and risk activities at tier 1 and tier 2. Activities at this level include categorizing the information system; implementing security controls, and managing the selection of the implemented controls, including continuous monitoring. The information system in this tier is at the core of the risk assessment process and dependent on the accurate implementation of security controls, including common controls across all three tiers in order to operate as effectively as possible. Controls not allocated to tier 1 or tier 2 are levied to the information system at tier 3.

INSIGHT Observations on the Risk Management of Medical Device and Software Cybersecurity

People don't do what we expect, but what we inspect.

Lou Gerstner
Former chief executive officer (CEO) of IBM

Many healthcare organizations have worked hard to reduce cybersecurity risks in recent years. Annual risk assessments have been completed, often by third-party security professionals. New technology has been implemented. Security-related processes have been improved. Additional staff with cybersecurity skills have been hired. Progress has been made. Unfortunately however, there is still much to be done before patients can have the level of confidence in the cybersecurity of healthcare organizations that they deserve.

One area of potential weakness is the supply chain and those vendors who provide applications and medical devices to healthcare organizations. Do they measure up to healthcare security requirements? It is not an easy question to answer, but it is an important one.

In recent years, the cybersecurity risks related to medical devices have captured much attention. Recently, a *Reuters* article stated, "The U.S. Department of Homeland Security is investigating about two dozen cases of suspected cybersecurity flaws in medical devices and hospital equipment that officials fear could be exploited by hackers." The devices under investigation include infusion pumps and implantable heart

devices. None of this is new. In 2012, Jerome "Jay" Radcliffe described how he made the troubling discovery that he could hack the insulin pump which he wears and on which his life depends. In addition to the risk that life-sustaining medical devices may be hacked, and their settings altered to cause injury, there is the risk that medical devices may become infected by malware or be attacked in some other manner that exploits a cyber-vulnerability. This may cause the medical device to malfunction or result in the device being used by cybercriminals for their nefarious purposes.

How are we to respond to these cybersecurity risks in medical devices?

1. Know what the Food and Drug Administration (FDA) requires of manufacturers and healthcare organizations with respect to cybersecurity:
   - "Manufacturers are responsible for remaining vigilant about identifying risks and hazards associated with their medical devices, including risks related to cybersecurity, and are responsible for putting appropriate mitigations in place to address patient safety and assure proper device performance. Hospitals and healthcare facilities should evaluate their network security and protect the hospital system."
2. Don't purchase medical devices that can store electronic protected health information (ePHI) without examining the Manufacturer's Disclosure Statement for Medical Device Security (MDS^2). This form discloses the cybersecurity safeguards supported or not supported by the device. This information can be used to evaluate whether a medical device has sufficient security safeguards. Do you, for example, want to purchase a medical device if the manufacturer states that they do not test and approve operating system security patches or allow antivirus software to be installed?
3. Ensure that medical devices are installed and operated as per organizational security policy. Important considerations may include vendor remote support processes, secure network configuration, installation and updates of antivirus software, installation of security patches as approved by the manufacturer, and back up of data.

Healthcare organizations need to be informed of the cybersecurity risks in medical devices and should make it clear to the manufacturers that purchases will not be made without evidence that cybersecurity risks are being adequately addressed.

Like medical devices, applications such as electronic health records (EHRs) and patient portals have similar risks. How do healthcare organizations know that an application is secure enough to protect ePHI? Cybersecurity weaknesses in applications have led to numerous security breaches. During risk assessments, hospitals should ask their application vendors for evidence of third-party security assessments of each application.

One recent response from a major healthcare IT vendor states that they have conducted an internal security risk assessment, the results of which are confidential, and that they did implement and verify risk control measures, but that there has not been any third-party security review. The vendor also states that the product has received meaningful use certification. Unfortunately, this is a pretty typical response. While it is reassuring to know that this vendor performed a risk assessment and implemented improved security measures as a result of the assessment, it is not too much to expect third-party validation of the application security and greater transparency about the results. The stakes are just too high to accept less.

If you are wondering if the EHR certification criteria dealing with security controls are sufficient, unfortunately the answer is no. While the application security controls addressed in the certification criteria are important, application security risks go far beyond what is addressed in the certification criteria.

The financial industry is ahead of the healthcare industry with respect to cybersecurity. The Financial Services Information Sharing and Analysis Center (FS-ISAC) recently published a white paper dealing with concerns about the security of software from third party service and product providers. The paper makes the following observations:

1. "It is the responsibility of the financial services industry to make software security requirements explicit rather than implicit.
2. If a vendor controls the development and build process, then they also are responsible for applying appropriate security controls."

> The paper from FS-ISAC recommends three security controls. While all three controls are relevant to healthcare, I will describe just one. It is that financial services industry members should require their software vendors to utilize a software security testing methodology known as binary static scanning and that the scans be conducted by a third-party. The scans can expose security vulnerabilities in software and can be used by the vendor to fix security vulnerabilities before the software is released. A scan report summary can be provided by the vendor to the financial organization to demonstrate that third-party security reviews have been conducted and to communicate the findings.
>
> The financial services industry has developed an effective approach to assure application security in its third-party developed applications. In healthcare, there is clearly a similar need since virtually all healthcare applications are developed and purchased from third parties. We need to communicate similar expectations. I encourage your feedback and suggestions on this important matter. (Bell 2014)

Managing ICT Security Risk through Governance, Control, and Audit

Earlier in the chapter you learned that confusion about managing specific roles within the organization to establish cybersecurity and risk management best practices frequently leads to compromise of information. Often, information will exist in two different areas of the organization and modifications made, in turn, lead to inconsistencies that make cybersecurity difficult if not impossible.

The best way to ensure that the compromise does not happen is to implement and sustain ICT governance program within the organization. Such a program should include strategies and policies that put technical and behavioral controls in place to safeguard all of the hardware and software that need to be protected. Security controls are technical or administrative safeguards or countermeasures for avoiding, counteracting, or minimizing loss or unavailability due to threats acting on their matching vulnerability, i.e., security risk. Controls are referenced all the time in security, but they are rarely defined. Further, mechanisms should be in place to ensure the employment and efficient use of controls throughout the ICT process. This can be achieved through performing audit functions that have been clearly defined in an audit plan. The purpose of this section is to take a detailed look

at governance, controls, and audits in order to further understand their role within the scope of risk management as a vital component of cybersecurity. The remaining chapters of this book look at each in terms of their roles within the NIST Framework for Improving Critical Infrastructure Cybersecurity and the COBIT 5 framework.

Governance

You have already read that protecting data is priority within most organizations in light of the recent problems at the National Security Agency, Target, and other organizations. Developing policies represents the first step for any effective risk management and compliance program. Where do we start? Policies help align the organization to management's vision, effectively communicating how leaders wish the organization to operate and providing important guidance to management. Most managers welcome these guidelines when determining their course of action. While many policies appear to be obvious, most organizations implement governance based on established frameworks (such as the NIST framework and COBIT 5). Frameworks help guide the development of a set of policies concerning a particular area of risk or compliance. Other examples include well-known documents such as the International Organization for Standardization (ISO) standards ISO 9000; the Project Management Institute's Project Management Body of Knowledge; and Committee of Sponsoring Organizations (COSO), the framework that supports all Sarbanes–Oxley compliance programs. Some may view these documents as standards, which is true, but policymakers use them as frameworks to develop a set of comprehensive policies.

Most organizations share similar risks and objectives. Industry groups establish frameworks to try to address the same types of concerns. ISO 9000 became popular when many companies were attempting to improve the quality of their processes and products. This framework represented a proven approach, and thousands of companies became certified as a result. Likewise, the COSO framework has become the internationally recognized standard for financial reporting.

Developing an effective set of policies is a top–down effort based on an established framework, sensitivity to the objectives of the

organization, the risks management faces, and regulatory compliance. Developing policies to protect sensitive data led us to several frameworks: We have already mentioned several NIST special publications. ISO 27001 has become the international de facto standard for information security management. Another widely used standard is the Payment Card Industry Data Security Standard. Since the early 2000s the management within the healthcare industry has had to become familiar with the Health Insurance Portability and Accountability Act, which provides privacy and confidentiality of all patient health data and information. A final example is Internal Revenue Service Publication 1075, *Tax Information Security Guidelines for Federal, State and Local Agencies*. All of these documents describe a series of recommended controls designed to address data security risks. Thankfully, they include the same types of controls and follow similar strategies.

The cornerstone of any effective risk management and compliance project is the risk assessment and strategic business plan. Policies support the organization's governance by meeting stakeholder needs and addressing risk. The upper-level management communicates their strategy through a collection of policies provided to the management. The risk assessment highlights the most important areas of concern and allows management to construct policies that address these areas, and most policies follow established frameworks.

The next thing we need to consider is implementation. Many ICT organization approaches borrow from change management strategies utilized during any significant organizational change. These strategies include education, management support, communicating the need to change and the benefits of the future state, and encouraging ownership of the new policies.

As is the case in all ICT strategy implementation, chances of success are much higher when the key members of management are involved during the entire process. Involving managers early helps them overcome fears about changes, and they view the new set of policies as partly theirs. Well-publicized security breaches can help them understand the overwhelming risks of not changing.

Documenting and implementing policies require significant effort and investment. Ensuring that the policies adequately address the organization's needs is also critical. Frameworks assure management

that it is on the right track and that the result will address the most critical areas of concern.

Controls

According to the Government Accountability Office, "The control environment sets the tone of an organization, influencing the control consciousness of its people. It is the foundation for all other components of internal control, providing discipline and structure. Control environment factors include the integrity, ethical values, and competence of the entity's people; management's philosophy and operating style; and the way management assigns authority and organizes and develops its people" (United States Government Accountability Office 2009).

Often, controls are grouped into one of two categories: general controls or application controls. General controls govern the design processes, security procedures, and use of the software and security of the data files throughout an organization's ICT infrastructure. From this we can derive that general controls apply to all computerized applications and consist of the combination of hardware, software, and manual procedures that create an overall control environment.

General controls include software controls, physical hardware controls, computer operation controls, data security controls, controls over implementation of system processes, and administrative controls. Table 1.1 describes the functions of each of these controls.

Application controls are unique to each individual application within the system, such as payroll or order processing. They include both automated and manual procedures that ensure only authorized data are processed by the application. Application controls are often further classified as follows:

- Input controls
- Processing controls
- Output controls

Input controls ensure data accuracy and completeness. Such controls include input authorization, data conversion, data editing, and error handling. Processing controls ensure the data completeness and accuracy during processing. Output controls ensure that the results of the processing are accurate, complete, and properly disseminated.

Table 1.1 General Controls

TYPE OF GENERAL CONTROL	DESCRIPTION
Software controls	Monitor the use of system software and provide unauthorized access to application and system software.
Hardware controls	Ensure that system hardware is physically secure, and check for equipment malfunction.
Computer operation controls	Oversee the work of the ICT organization or function to ensure that defined procedures are consistently and correctly applied to the storage and processing of data. They include controls over the setup of batch processing, as well as backup and recovery procedures for processing that ends abnormally.
Data security controls	Ensure that business asset data on disk or tape are not subject to unauthorized access, change, or destruction while they are in use or in storage.
Implementation controls	Audit the ICT development process according to an adopted framework, to ensure that the process is properly controlled and managed.
Administrative controls	Formalize standards, frameworks, rules, procedures, and control disciplines to ensure that general and application controls are properly executed and enforced.

Notice the correlation between application control categories and the generic definition of a computer as being any electronic device designed to accept input, perform processing, and produce output. The generic definition typically also includes a feedback loop. In this context the feedback to the organization regarding appropriate design and implementation of the controls comes from audits, which are discussed in the next section.

Audits

Putting controls in place based on a framework supported by strategies and policies defined through an ICT governance program is a good first step toward fulfilling the requisites of cybersecurity risk management. However, management needs to know that ICT security and controls are effective. This is achieved through conducting complete and systematic audits. An ICT audit examines the organizations over all security environments in addition to the controls that govern individual information systems. COBIT 5 provides a comprehensive set of audit process definitions that an organization can implement that allows an ICT auditor to trace the flow of sample transactions

through the system and perform tests using, if appropriate, automated audit software. The ICT audit may also examine data quality.

Security audits review technologies, procedures, documentation, training, and personnel. A thorough audit plan is created and may also include a simulated attack or disaster to test the response of technology, ICT staff, and individuals from other business units.

The audit lists and ranks the control weaknesses and estimates the probability of their occurrence. Next, it performs an assessment of the financial and organizational impact of each threat. This information can then be used as feedback into the process, for management to be advised of the weaknesses and enable them to respond through a planning effort devised to counter significant weaknesses in controls.

INSIGHT A Walk through 3 Stages of an SAP Security Audit

Tracy Levine, a Systems, Applications, and Products in data processing (SAP) application consultant at itelligence, fields some questions about various stages of preparing for an SAP security audit.

In her blog post, "How to Survive an SAP Security Audit," Tracy Levine, an SAP application consultant at itelligence, writes about three stages of an SAP security audit and uses political terms to describe them: state of the union, political reform, and ongoing legislation. Using her model, I asked Tracy to respond to a few questions related to preparing for an SAP audit.

In the first stage, the state of the union, an organization asks itself the following questions: "Where are we now? How did we get here? What challenges do we face?" Could you cite some examples of challenges organizations may face when preparing for an SAP security audit?

The primary concern when preparing for an SAP audit is gaining a clear understanding of client-specific security requirements and being able to articulate these requirements in terms of business processes. Oftentimes, SAP security requirements remain undocumented, what we refer to as "tribal knowledge," as in there is no central repository that clearly defines the SAP landscape and tracks and monitors changes. As the SAP landscape matures, scalability efforts prove more difficult with increased requirements and a lack of correspondence between functional silos. Changes done by one functional team may override requirements that were purposely implemented for another, driven by a lack

of visibility with regard to change management. Furthermore, clients may be concerned with underlying segregation of duties (SoD) violations and the need for a Sarbanes–Oxley (SOX)–compliant deployable role design.

In his video from Governance, Risk and Compliance (GRC) *2013, Steve Biskie refers to thinking about risk in terms of "what can go wrong." He uses an example of "inadequate mapping of business processes to role design" as an example of something that can go wrong. Can you cite some other examples of what can go wrong in relation to managing the security of an SAP environment?*

To piggyback on Steve Biskie, suboptimal role designs, which are not task based, may lead to the provisioning of roles with too much access or inherent SoD violations. By utilizing a task-based role design, organizations are more inclined to adhere to the principle of least privilege, defined as a system in which users are only able to access the information and resources that are necessary for legitimate business purposes. Position-based or user-based role designs do not take into account scalability for business solutions and the opportunity for avoidable SoD conflicts.

Another example of mismanagement of security in the SAP environment involves inefficiencies surrounding the user provisioning process. A lack of automation with workflow capabilities and an embedded risk analysis can decrease visibility, and increase the risk of SoD conflicts. Furthermore, manual provisioning processes can lead to long cycle times from the time of request to the time access is granted or denied.

How has the emergence of mobile and cloud technologies posed new challenges to organizations that are preparing for an SAP security audit?

With the expansion of mobile and cloud technologies comes an increased risk for cyber attacks to SAP systems. This leads us to the question, how secure is the cloud and is there a way to mitigate any associated risks? One of the benefits of the SAP HANA Cloud Platform, as highlighted at the 2013 SAPPHIRE Now conference, is the ability to leverage multiple deployment options—whether in a customer's data center, the public cloud, the managed cloud, or a hybrid environment—to help meet the changing needs of any organization. SAP has also released SAP Mobile Secure, an enterprise mobility management tool that provides organizations with increased security for apps and mobile devices. Furthermore, the cloud edition of SAP Afaria addresses the need for a convenient, reliable, and low-cost solution that provides a way to manage security risk with or without any prior SAP infrastructure.

In the second stage, political reform, one of the questions you cite that an organization asks is "How will this change the way we do business?" Could you provide an example of a business-changing issue pertaining to security of an SAP environment?

Political reform refers to the need to make modifications to the current role design, role management methodologies, or user provisioning process. This can lead to a need for clearer definition of jobs and responsibilities in the SAP landscape. Additionally, it may require previously grouped tasks to be separated across individuals within an organization to avoid inherent SoD conflicts. For some companies, the greatest challenge in prioritizing SAP security is the need for increased collaboration between the business and IT. This collaboration can lead to more defined requirements regarding ownership of SAP roles and critical permissions across functional areas within the organization.

In the third stage, ongoing legislation, one of your questions is "What risks do we still have and how are we going to monitor them?" Could you provide an example of changes an organization has made to the methods it uses to monitor risk?

There are many automated and manual options when it comes to monitoring and assessing risks. SAP Access Control and Risk Management offer tools to meet this demand. The Risk Management tool enables organizations to quantify risks based on impact and probability. Owners can be assigned to risks, and notifications can be activated to bring awareness to controls that have or have not been performed. Those who are fearful of an upcoming SAP audit need not worry. Most organizations have unavoidable risks associated with some aspect of their security design. However, it is essential that companies are able to demonstrate the steps they have taken to mitigate these risks. Leveraging automated or manual mitigating controls as part of SAP Access Control is one method for taking a proactive approach to known risks.

For organizations that do not deploy GRC, it is essential to have a central repository to manage approvals and changes that have been executed in the SAP landscape. A repository can be used as an in-house audit tool to track change logs or actions that have been taken against mitigating controls. However, risk-monitoring methods are only valuable if companies have been able to appropriately assess not only risk priority levels but also the effectiveness of controls that are in place.

When you speak with clients about preparing for an SAP security audit, what do they seem most concerned about? What's turning them into insomniacs?

First of all, no one should be losing sleep over an SAP security audit, but a lack of transparency can lead to uncertainties surrounding the business. The greatest concern, understandably, is a need for increased visibility into the cumulative nature of a user's access in business terms. This lack of visibility also leads to risks with regard to critical access and permissions. Many organizations are concerned that they aren't following industry standards and best practices when it comes to SAP security, which will evoke wariness in auditors. Has everything been secured as it should be? Are we following a streamlined process when it comes to making critical program and configuration changes and have the necessary authorization checks been implemented and tested? Who requested the change, who approved it, and who tested it? This again is where the need for a central repository or GRC landscape comes to fruition. It is not enough to have the processes in place. Monitoring efforts need to be performed to ensure that controls have been executed.

Another growing concern for savvy organizations is the need for information in real time. Many companies have employed detective controls, which do not allow for a proactive approach to SAP security risks. Detective controls are valuable with regard to reporting and analytics, but with the onset of competition in the environment via cloud and mobility comes an increased need for preventive measures and risk-monitoring efforts. By gaining a better understanding of the organization's SAP landscape, organizations often see decreased administrative and testing efforts to maintain the security environment, an increased ease of scalability, and a minimization of SoD conflicts as well as data integrity issues due to a misuse of users' authorizations. (Byrne 2013)

Implementing Best Practices Using a Single Cybersecurity Framework

Formal standards are meant to embody the model for the "common body of knowledge and accepted state of industry best practice" (NIST Framework for Improving Critical Infrastructure Cybersecurity and COBIT 5). Logic should support the correctness of that assumption without much additional proof, as it would be impossible to build a common cybersecurity governance, risk management, and audit

process without adopting some sort of standard, accepted model. Given the amount of activities embodied in ICT work, such a model has to be broad and comprehensive. As a result, such all-embracing models are commonly called umbrella frameworks.

Umbrella frameworks are named after their intent, which is to cover the entire scope of ICT that they define. In that respect, umbrella frameworks specify an ideal model at a level sufficient to allow any organization to tailor processes to fit within its structure. Theoretically, a single umbrella framework can describe a competent technology or management process in any level of detail.

Many of the umbrella frameworks utilized in the ICT industry provide specifications for the activities performed in software and system life cycle processes. Noticeably absent in those frameworks are specifications for activities related to governance, risk management, and audit of cybersecurity processes. You will learn in the next several chapters that, as much as there is a life cycle for activities performed to develop ICT systems and software, there exists an umbrella framework providing specifications for cybersecurity life cycle processes. Further, you will learn that many of the defined cybersecurity activities can and should be performed in parallel with the activities of the other industry frameworks.

Nonetheless, you should keep two important caveats in mind when considering the application of umbrella frameworks. First, no two organizations operate in the same way, so each individual cybersecurity life cycle process has to be considered differently in terms of its particulars. These differences may be large or small, but because they exist, every organization must decide how to explicitly array the processes it adopts within the larger concepts and principles of a cybersecurity life cycle model represented by an umbrella framework. In other words, although organizations can use a standardized framework to guide the creation of a coherent set of defined governance, risk management, and audit processes and activities, they must tailor the implementation in a way that makes the most sense for them.

Second, an organization should not rely on just the process definitions of the umbrella framework. All umbrella frameworks reference, in one way or another, other industry standards to further define the activities and documentation required of each life cycle process. Our discussion of the Framework for Improving Critical Infrastructure Cybersecurity

is no different. That framework provides the foundation for the definition of cybersecurity life cycle processes, and relies on COBIT 5; the Council on CyberSecurity's Top 20 Critical Security Controls; ANSI/ISA-62443-2-1 (99.02.01-2009), Security for Industrial Automation and Control Systems: Establishing an Industrial Automation and Control Systems Security Program; ANSI/ISA-62443-3-3 (99.03.03-2013), Security for Industrial Automation and Control Systems: System Security Requirements and Security Levels; the ISO/International Electrotechnical Commission (IEC) standard ISO/IEC 27001, Information Technology—Security Techniques—Information Security Management Systems—Requirements; and NIST Special Publication 800-53 Revision 4: Security and Privacy Controls for Federal Information Systems and Organizations to further define each activity.

Chapter Summary

- An organization's digital data are vulnerable to destruction, misuse, error, fraud, and hardware or software failures. Managers must be aware of which data and information assets are most vulnerable and valuable in order to implement the appropriate safeguards.
- Lack of strategies and policies developed through an adopted cybersecurity governance program, and controls that support those strategies and policies, can cause organizations relying on business functions to lose sales and productivity. Data and information assets, such as confidential employee records, trade secrets, or business plans, lose much of their value if they are revealed to outsiders or they expose the organization to legal liability.
- Organizations need to establish a good set of general and application controls for their information infrastructure. Processes associated with risk management are defined to evaluate information assets, identify control weaknesses, and determine the most cost-effective set of controls.
- The ICT cybersecurity life cycle is composed of a coherent set of best practices based on a framework that defines the activities and tasks that should be performed.

Case Project

On Christmas Eve 2014, an unwanted Christmas present was delivered to Suny Corp. That afternoon Suny's network was attacked, attracting numerous phone calls from customers angered at an error message they were receiving when they attempted to access the network. An independent hacker group called "Lizard Squad" took responsibility of that attack and claimed that it was in retaliation of a recently released movie. This attack has generated concern by Suny management about the security of not only their networks but all of their data and information assets. They would like an assessment of their existing risk management procedures and ask that the NIST Framework for Improving Critical Infrastructure Cybersecurity be used to redefine their cybersecurity life cycle process.

As a first step, your task is to help Suny assess their existing governance, controls, and audit procedures. Since this situation is real, but you do not have access to the internal management and existing ICT infrastructures, your results will be hypothetical. However, you should provide a statement detailing what exists currently, so that you can accurately develop a plan for the implementation of the framework once you become better acquainted with risk management frameworks in the next chapter.

CHAPTER

2

ORGANIZATIONAL RISK FACTORS FOR UNINTENDED INSIDER THREAT

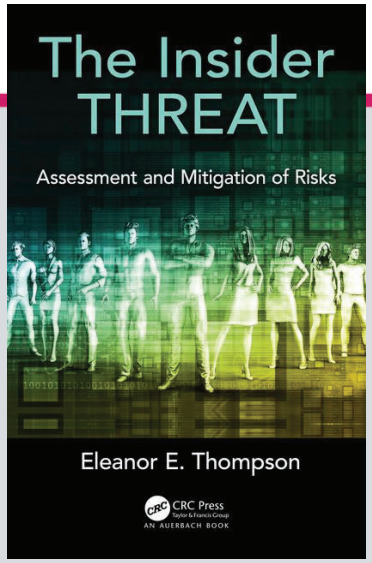

This chapter is excerpted from

The Insider Threat

by Eleanor E. Thompson

Learn more

3

Organizational Risk Factors for Unintended Insider Threat

3.1 Introduction

An ideological shift must be made in terms of insider threat so that more aggressive and more thoughtful mitigation techniques can be applied in a more holistic manner. This particular chapter addresses risk factors that are not commonly understood, nor likely/frequently considered as part of the insider threat program. The risk factors that contribute directly to unintended insider threat is where a great deal of progress can be made to assess and mitigate risk within organizations. If leadership considers and implements these categories as part of a risk mitigation program, then a better change process along with a greater understanding of their associated mental models, will occur.

A greater depth of understanding is needed for organizations to more comprehensively understand the risk factors that reside within their organizations. How the business work function is conducted today looks significantly differently than it did a generation ago, including outsourcing of work. The option for a significant number of employees or contractors to extend their work environment into the vast world, seemingly anywhere, is now common practice, well beyond the traditional brick-and-mortar walls owned by an organization. Contracts and subcontracts create opportunities in which tasks performed may be far beyond the observation of traditional security protocols where opportunities for response may or may not be immediate. When operating system updates or patch updates are delayed, vulnerabilities may exist when computers are not properly managed or tracked as an asset. Unsupported and out-of-date software, where there are no vulnerability patches, can create havoc very quickly in organizations. Additionally, some work spaces alone may be in higher

threat locations such as: at home, on a metro, in a coffee shop, at a remote telework site; hotel bars and lobbies, and other contractor designated work spaces. Globalization has created vast amounts of business that is being conducted on-the-go. The places from where network connections can be established, given the right equipment and passcodes, are seemingly endless. Human behaviors are difficult to monitor by security personnel in such disparate locations, if at all. It is unlikely that this type of remote work will decrease due to increasing population centers. Employers have sought to decrease both their transportation and corporate office carbon footprints by requiring their employees to telework and in some cases offer telework as a benefit in an attempt to retain quality employees.

First, I provide a narrative response, as I did in the preceding chapter, that directly answers what factors contribute to unintended insider threat, then I provide a listing of categories that are displayed in Table 3.1. These categories are listed for readability in terms of core categories as well as their cascading sub-categories. Last, toward the end of the chapter, I enter into a more comprehensive discussion. Throughout the discussion, practical application, mental models and best practices are called out. Specifically, the categories may be used by an organization to identify controls that they should use to mitigate the risk element.

What are the factors that contribute to unintended insider threat? The answers follow.

3.1.1 The Narrative Response

- Unintended insider threat increases with human risky behavior; this behavior is captured into typologies that include the following employee behaviors: transporter carrier; litterer; too-tolerant being; risk revealer; scattered data collector; seldom unexpected; gluttonous optimizer; unknowing curator; defiant disregarder; technically obsolescent; hygiene-hindered communicator; goofing around player; blissful user; trust believer; rushed surge responder; memorable password scribe; privileged user abuser; unmonitored teleworkers; wide-open wanderer; verbal space cadet; unguarded doorman; absentminded; face-off subordinate; shortcut alleyway taker; uninstructed newbie;

surprisingly superhuman; ivory towered; road traveler; part-present, part-timer; storyteller, too-busy-to-tell; not-checked-out employee; the laissez-faire trainee.

- Physical environmental risk factors, coupled with aggravated risk mitigation approaches, and gaps in architectural wellness can lead to increased unintended insider threat risk. Ensuring well-coordinated efforts between physical and cybersecurity mechanisms, including incident response is important. Minimizing the unknown factors among various units/staffs should be a goal. Recognizing that systems may be designed or propagated with error and coordinating to ensure that future changes minimize exposure should also be a goal to minimize unintended insider threat.

3.1.2 The Categorical Response

Now that the narrative has been presented, I will now list the categorical results. These categorical results may be used in terms of an assessment, to identify the organizational responsibility or lead, and assigned control. Following the table is a more descriptive discussion of each of the Core Categories.

Table 3.1 Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
Core	**A. Human Behavioral Risk Factors Influence or Contribute to the Outcome of Unintended Insider Threat**	**Suggested type of insider threat with identified typology**
s	*Typology 1, the transporter – carrier.* Those who carry data back and forth are more at risk for losing it (transporters, more easily lost, or unattended).	Virtuous, Wicked
s	*Typology 2, the litterer.* People who throw out data or leave things unattended.	Virtuous, Wicked, Vengeful, Malicious
s	*Typology 3, the too-tolerant being.* Humans, such as managers and co-workers allowing others do things they shouldn't.	Virtuous, Wicked, Vengeful
s	*Typology 4, the risk revealer.* Being tricked, social engineered, shoulder surfed. Examples include: intentional disclosure, or unintentional, accidental, office function disclosure such as operational information inappropriately released, writing information that can be compiled into a pattern for gain or other exploitation. Others may be waiting for this person to make a mistake.	Virtuous, Wicked, Vengeful, (Malicious)

(*Continued*)

Table 3.1 (*Continued*) Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
S	*Typology 5, the scattered data collector.* Collecting information without protection, supply chain consideration, end user, potentially breaking the chain of custody, disruption.	Virtuous, Wicked, (Vengeful)
S	*Typology 6, the seldom unexpected*. Those who do not alter their behavior and are predictable, allowing for others to be familiar.	Virtuous
S	*Typology 7, the gluttonous optimizer*. Seeking gain beyond what is authorized or normal behavior in the work environment may also be convenience motivated.	Wicked
S	*Typology 8, the unknowing curator.* The holder of treasure, or very important things, but unaware of its importance or consequence if compromised. This is an example of someone not aware of immediate or further cascading impacts, should treasure be compromised.	Virtuous
S	*Typology 9, the defiant disregarder.* Ignoring known security practices, or training.	Wicked, Vengeful
S	*Typology 10, the technically obsolescent.* Various levels of this, individuals' technical lack of knowledge or skill to protect data, or unknown reasons.	Virtuous, Wicked
S	*Typology 11, the hygiene-hindered communicator.* Poor receiving and sending hygiene; for example, sending to an unconfirmed agency, person, office, individual, before verifying legitimacy of receiver. May send unprotected when password protection is needed.	Virtuous, Wicked
S	*Typology 12, the goofing around player.* Someone who just plays around to see what they can do, or someone who plays jokes.	Virtuous, Wicked
S	*Typology 13, the blissful user.* Basic and privileged users that have a lack of experience or training.	Virtuous
S	*Typology 14, the trust believer.* Those who are more prone to being socially engineered, or unguarded. Keeping in mind that most everyone can be tricked or a social engineering target. A social engineering targeted attack, also known as spear phishing.	Virtuous
S	*Typology 15, the rushed surge responder.* Missing information assurance factor. Information assurance may decrease during times of immediate incident response, leaving gaps between traditional/physical security and cybersecurity. The needs to provide information may prevail beyond ensuring best security practice will be hierarchical, or externally politically driven.	Wicked

(*Continued*)

Table 3.1 (*Continued*) Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
S	*Typology 16, the memorable password scribe.* The user who cannot remember numerous passwords between home and work, and writes them down, and may transfer back and forth, either online or in person. Collects work and home passwords together.	Virtuous, Wicked
S	*Typology 17, the privileged user abuser.* The ultimate MacGyver work-around with permissions. May create work-around for own benefit to include pornography, gaming, or other interests.	Wicked
S	*Typology 18, the unmonitored teleworker.* This goes beyond work productivity oversight to information security hygiene at home.	Virtuous, Wicked, Vengeful, Malicious (not monitored, could be any)
S	*Typology 19, the wide-open wanderer.* Leaves terminal open or provides access unintentionally or willingly.	Virtuous, Wicked
S	*Typology 20, the verbal space cadet.* In-their-own-world talkers. Those who are on the phone and they ignore physical surroundings and/or are less aware of surroundings and speak too loudly about information that should be protected. This is becoming more common place. Changes in physical space could amplify this—office vs. cubical. Could be seen as loud mouth or overzealous.	Virtuous, Wicked
S	*Typology 21, the unguarded doorman.* More easily allows intrusions, does not question legitimacy of those present physically, nor recognize virtual open doors.	Wicked
S	*Typology 22, the absentminded.* Those who forget small details, not gross negligence.	Virtuous
S	*Typology 23, the face-off subordinate.* Those who dislike their supervisors and do not wilfully make things easy for their supervisor, nor look out for their best interest; outright sabotage, or withholding of vital information for example. The subordinate may still be true to their perception of ethics, and love for the organization.	Vengeful
S	*Typology 24, the shortcut alleyway takers*. Those who want things to be on Easy Street and ignore training and policy.	Wicked
S	*Typology 25, the uninstructed newbie.* Those who are unfamiliar with a particular system. They may not be technologically obsolete but have not had enough experience on a particular system and are not familiar with all of the processes, functions and capabilities of the system including risk in their actions.	Virtuous

(*Continued*)

Table 3.1 (*Continued*) Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
s	*Typology 26, the surprisingly superhuman.* The individual who can hear or see information that they are not supposed to without seeking out this information because of others' behaviors such as talking too loudly.	Virtuous, Wicked, Vengeful, Malicious
s	*Typology 27, the ivory towered.* The user that confuses position status, or other social or ranked privileges, with good information technology hygiene, and may ignore policy. Senior positioned folks (managers, executives, for example) not setting the example, and not adhering to policy.	Wicked
s	*Typology 28, the road traveler.* Using very unsecure networks to access systems, including hotel networks without understanding the associated risks, and not taking precautions to turn off higher risk applications.	Virtuous, Wicked
s	*Typology 29, the part-present part-timer.* Those who are not always present, and don't consider their job primary. Short cuts more frequently taken to access or share data. Data may be repeated on servers throughout the world.	Wicked
s	*Typology 30, the storyteller, too-busy-to-tell.* Storytelling plays a key role in learning organizations. Cybersecurity needs more organizational and timely stories to translate into an analog world in order for more organizational resilience. Lessons learned may be considered part of a story telling.	Virtuous, Wicked
s	*Typology 31, the not-checked-out employee.* Departs job function or status by leaving one position to go to another that requires different types of access. For example, as military separate or retire, then check in as a civilian or, a contractor. Because of transition, they could retain system privileges, authorities, distribution lists in lieu of cancellation. Privileged users must be checked out properly from any organization.	Virtuous, Wicked
s	*Typology 32, the laissez-faire trainee.* Bypasses training, doesn't keep up with system-specific training. Doesn't pay attention in training.	Virtuous, Wicked
Core	**B. Organizational Process Risk Factors That Influence or Contribute to the Outcome of Unintended Insider Threat**	**Lead/control (Use as a checkoff sheet to determine if and where this fits in your organization)**
s	*Audit risk factor.* Audit systems should be in place in an organization. This includes conducting physical and cyber audits on a random basis, as well as routine.	

(*Continued*)

Table 3.1 (*Continued*) Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
s	*Sending outside organization's network.* Several risks may be mapped to this including unconfirmed addresses, office, people, proprietary nature of information sent, restrictions on that information, violating protocols.	
s	*Policy risk.* As part of organizational process, policy plays a role in establishing protocols. However, policy may not be followed, or can be perceived as being poor policy and therefore operators may act differently. Policy may instead create perceived roadblocks to inhibit practical operations and may not be enforced organizationally. Additional policy could have other consequences counter to intention.	
s	*Money.* Lack of organizational prioritization or understanding of acceptable risk to adequately prioritize specifically known security risk mitigation measures.	
s	*Training and education.* Training and education may or may not be in place. Training and/or education may not be current enough or adequate for continually emerging changes.	
s	*Chain of custody risk.* Not understanding sensitivity of data in the chain of custody of that data extends from collection to end state of data.	
s	*Management of personnel behavior risk.* Not understanding red flags in personality, disorders, behaviors, and/or changed behaviors in staff.	
s	*Communications risk to senior leadership.* Staffs not being able to present to senior leadership in the right way, or the rules of the cyber world, to restrict it.	
s	*Organizational reputation.* Making the organization look bad, may be internal or external.	
Core	**C. Physical Environmental Risk Factors**	**Lead/control (Use as a checklist to determine if and where this fits in your organization)**
s	*Unprotected equipment.* Physical barriers not always in place or properly used to protect equipment. Especially check locks, and who has the key.	
s	*Internal physical exposure to different agencies and contractors.* Allowing knowledge that extends beyond need-to-know due to physical proximity of work spaces.	
s	*Computer security access.* On-screen access should be limited to authorized viewers/users. Positioning of monitors of cubical constructions.	

(*Continued*)

Table 3.1 (*Continued*) Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
s	*Illusion of privacy and security.* Illusion of privacy while talking on phone or using other technology.	
s	*Mask of familiar interface.* Familiarity does not mean safe and secure.	
s	*Lack of cyber-physical digital analog bilingualism.* Not understanding how one or the other is different. For example, mailing an email, not the same as sending letter home. Email has extensive repeatability on servers around the world.	
s	*Lack of equipment accountability.* Not keeping proper track of equipment.	
s	*Physical geographic misperceptions.* Thinking in terms of analog geography vice cyber reality; threat doesn't depend on distance.	
Core	**D. Architectural Information Technology System Wellness**	**Lead/control (Use as a checklist to determine if and where this fits in your organization)**
s	*Openness.* Open ports and file transfer protocol. System architecture that allows for open discovery, and not locked down. This means it is not restricted.	
s	*Privileged user.* More access to systems and granted greater permissions within the system to utilize or change the system.	
s	*A basic user with root access.* A user who does not understand high-level permissions and could delete valuable system data and enterprise files unknowingly.	
s	*Role-based access.* System may allow for greater access than is necessary, and unless role-based access is implemented too much access may be granted.	
s	*System-of-sending.* Unsecure, could be visible to others if protections are not used.	
s	*Error propagation through technological facilitation.* An error being repeated due to system set ups, such as reply all or distribution lists that may contain contacts that should not be used because of less secure email addresses.	
s	*Seamless technology.* Technology can appear to be too seamless to understand the risk. People may become too comfortable with systems that are linked together, such as phones, and phone applications with various security levels.	
s	*Enterprise change management.* If not properly engineered, changes could expose the network.	

(*Continued*)

Table 3.1 (*Continued*) Factors That Contribute to Unintended Insider Threat

CATEGORY	CORE CATEGORY = C, OR SUB-CATEGORY = S	LEAD/CONTROL
Core	**E. Aggravated Risk Mitigation Approach**	**Lead/control (Use as a checklist to determine if and where this fits in your organization)**
s	*Inappropriate or untimely action without needed collaboration.* Rushing to fix a perceived threat could have greater consequence than the threat itself. Coordination may be necessary before taking a system or person offline or performing a mass denial subprogram.	
s	*Mitigating cyber threat without reaching out to physical security management.* Those persons removed from system access could still pose a physical threat; coordination should be planned and required.	
s	*Moving too slowly.* Leaving identified risks open too long to fix a known threat (a balance is needed). Network evaluations may delay fixes.	
Core	**F. The Unknown Factors**	**Lead/control (Use as a checklist to determine if and where this fits in your organization)**
s	*Unknown consequence linkages.* Organizations may not be able to correlate from a single user or other event compromise to a particular outcome or consequence.	
s	*Unknown, but known threats.* Users or management may not be given latest threat information due to a need-to-know basis, or risk of leaking information regarding this vulnerability. Direction to avoid specific risk may be given, but trust is required to comply.	
s	*Unknown disclosures from business.* Disclosures of information may occur from other business sources for individuals or organizations. Difficult to track this without proprietary knowledge planning.	
s	*External storage.* Unknown how many people are storing hard copy or data work products where they should not.	
s	*Already known, but not shared.* Threat may remain unknown to an organization or agency, even though it may have been identified elsewhere.	

Insider threat is a phenomenon that can, in part, be observed empirically with a keen eye, although not always easily observed if people aren't trained to understand what type of behavior they should be seeking. Employees should understand what type of behavior is acceptable and that which is not acceptable. When incidents do arise, they should learn from their best practices. Indeed, individual motives and actions can and should be better understood that mitigation controls can be implemented. Primary examples of factors that contribute to unintended insider threat are presented as follows:

- Human behavioral risk factors as typologies
- Organizational process risk factors
- Physical environmental risk factors
- Architectural IT system wellness risk factors
- Aggravated risk mitigation approach
- Unknown factors

There are several themes of these factors to also consider. Themes that emerged have centered on *human behavior, organizational processes*, and the *physical environment*. Notably, actions that are intended to assist and mitigate risk could also cause additional damage and increase risk and could cause a cascading impact of unintended incidents. A well-coordinated response is a necessity for an organization to minimize additional risk. This includes communication between hierarchical levels, as well as between both traditional and cybersecurity offices. Unknowns, or unshared known information, are also emergent themes as factors contributing to unintended insider threat. Human behavior is a significant contributing factor for increasing risk within an organization. Risky human behavior extends well into the physical domain, beyond the user keyboard. Human behavior has emerged to become the weakest link even while responding to discovered threats. Not only that, but this behavior extends throughout the organization at all levels.

Along with human behavior, the organizational processes that humans have created, or failed to create, in conjunction with the physical environmental risk factors are of a concern. How well a broader architectural computer system is initially designed or subsequently modified is a risk factor that also includes particular assigned users, the privileged, the root based, and those with role-based access.

Even when employees try to do the right thing, such as identify a threat, and take that threat off-line they can still create a situation even worse than the initial threat.

Unknowns are often not just unknowns in the construct of general knowledge, but instead the distribution of that knowledge can be known to some parts of an organization and to a lesser degree in other parts of an organization. This demonstrates a greater need for information sharing and in some cases developing the processes, as well as accountability regimes to be more effective. The next section presents further discussion on the complete findings presented in Table 3.1.

3.2 Human Behavioral Risk Factors as Typologies

The cultural environment of an organization can help us better understand insider threat and influence and/or intervene on how individuals behave (or could potentially behave) in a social psychological context, within the groups that they belong, or interaction within the organization.

Although one mitigation strategy is to identify anomalies in employee behavior in the workplace: such as staying late; downloading or sending large files; or accessing information that is outside of an employee's scope of work. There are many other behaviors contained in the typology listings that can be equally monitored and managed to reduce risk against the organization today. Leads for managing this effort will need to be established. Typically, monitoring is something that an employee needs to be informed about, and that use of organizational equipment allows for detailed monitoring as well. However, monitoring of other behaviors mean that leads responsible need to collaborate on the best practice for this monitoring, as well as ensuring that civil rights and liberties are adhered to.

Typology 1, the *transporter-carrier*. Everyone loses things. I bet you have lost, during your lifetime, your car keys, your purse, your man bag, your credit cards, your orthodontic retainer, and yes—even your identification cards. At some point we have all misplaced or lost something. Some of these lost items go into a mystical land possibly where all the missing socks have gone; you know the land. However, some missing items are very much a threat and they do not simply vanish. These threats increase your risk. If you are a transporter of

work documents, either an electronic storage device or hardcopy form, you are living on borrowed time. It is all too common a scenario, placing a bag on the seat beside us while riding metro and the bag gets left behind, or quickly getting out of a hired car and again your bag is left behind. When an employee transports documents to and from work, they are at risk of losing information. Important information. It happens all the time, so don't do it. Despite well-meaning intentions, it only takes the briefest of moments or inattentiveness to lose an employer's intellectual capital or a client's personal information, your life will be forever changed. Common are reports of theft of personal information or proprietary information from unlocked cars, even in the driveways of their owner. Being a transporter or carrier is an increased human behavioral risk. Creating policy that limits what is allowed to be transported is a way to mitigate or minimize this risk. Having random spot-checks of employees leaving the organization, and oversight by those who know what can be removed and what cannot be removed will reinforce policy and help mitigate this risk.

If you are curious, I know I was, the *Uber Lost & Found Index* provides some interesting data on the items most commonly forgotten in their transports.[1] Notably, the most recent index demonstrates that commuters in the morning are less likely to forget, but that the spikes of forgetfulness of commuters increases between the hours of 4:00 p.m. and 7:00 p.m.; among the top 10 items of forgetfulness include bags and backpacks, which in my opinion could easily hold proprietary work products.

Typology 2, the *litterer*. This one is complex and requires extensive inspection of on-site security for periodic inspection of materials. Recycling materials can be attributed to the virtuous person who means well, but may recycle proprietary information that should be crosscut shredded or burned. Those who are leaving out information and/or stepping away temporarily can be monitored through traditional security programs such as unannounced walk-throughs and spot-checks on a random and frequent basis. It is quite difficult to identify those who are intentionally throwing out data to be vengeful or malicious. Having a secure electronic backup, and in some cases, hard copy documents in a secure location will minimize this loss should a need exist to retrieve them. However, ensuring that the supply chain of your trash is secure is equally important.

Spend the time to ensure that trash and recycling bins are periodically and randomly inspected. Ensure that the recycling contract contains provisions that direct periodic inspections of materials and verify that the contractor is not owned or connected to competitors in any way. Your company's security personnel should walk through the entire supply chain of waste management. If there are disposal bins for proprietary information, those bins should also be locked or kept in a secure location and not open for anyone to pick through or pick up sensitive information.

Typology 3, the *too-tolerant being*. A lack of accountability of peers, subordinates, supervisors, contractors, and others leads to this typology risk. An effective policy is key to mitigating the risk of individuals being too tolerant of each others' poor cyber hygiene practices. Of course, there tends to be some discretion left to supervisors, however violations that place the organization at risk, including human behavior, should be addressed and repeat offenders held accountable. Education and communication of policy throughout the organization is also effective in mitigating this risk. Accountability of personnel actions that run contrary to policy should be published in the annual report and used as a deterrent and a constant warning to staff of unacceptable behavior. This also sends a message to the workforce of the importance to adhere to policy.

Typology 4, the *risk revealer*. This risk is an example of those employees who are susceptible to disclosure, being tricked, socially engineered, or shoulder surfed. Almost all of us are susceptible to a certain level of social engineering and remaining alert of new risk areas is increasingly important. Work environments should have a robust social engineering prevention program, which includes training as well as periodic testing of employees' awareness. Pushing out the latest risks to employees can also keep vigilance up. Social engineering has become more specifically targeted, creating opportunities for exploitation over a longer period of time.

Loss associated with social engineering is no joke and can be painfully felt at multiple levels throughout an organization, down to an individual level. The more a threat (perpetrator) knows about an individual, the greater the potential to obtain access into organizations or banking institutions that the unknowing individual uses. Staying on top of the latest scams is important because many of these scams

can appear very legitimate. Fools are not the ones being socially engineered, the attempts are on all of us.

Typology 5, the *scattered data collector*. This is the employee that collects information but does not properly password protect it through encryption or password protection. This type of employee may reply to all or send to an email address that is not verified, or pulls the wrong distribution list onto their send list. The supply chain considerations from the collection to the movement and transfer to the end user is not properly secured. If you think in terms of evidentiary process, the chain of custody is disrupted. The information could then go on to be used for other means. Consider if you have ever sent or received an email that was not intended for you. Consider if it was a one-to-one case or if you were on a massive distribution list, which could not be recalled. Even if it can be recalled, within seconds a misguided email may be read, forwarded, and is now repeatable throughout the world. Mistakes can be made. When it is a very important email containing proprietary or private personal information, this can be especially devastating and the risk of liability or loss very realistic and sometimes not quantifiable for years.

Typology 6, the *seldom unexpected*. Individuals who do not alter their behavior and are predictable allow others to be familiar. This is part of basic security knowledge. The more an individual's pattern remains the same, the more predictable their behavior is to those who wish to compromise or infiltrate that pattern. Altering routines so they are not completely predictable may deter, in part, efforts to compromise an individual or area.

An example of how technology can take this to a high level of risk: In November 2017, a global heat map containing accumulated data points from the preceding two years of fitness trackers was published on the internet to highlight usage of the GPS technology. Whereas city centers showed mostly solid light, this was not the case in more remote areas around the globe where fitness trackers were being used by U.S. military personnel during routine physical fitness activity. For example, the posting of this global heat map, a conglomerate of information, remained unnoticed by the military community until Australian student Nathan Ruser[2] stumbled upon it. Ruser tweeted the following statement: "It looks very pretty, but not amazing for Op-Sec. U.S. bases are clearly identifiable and mappable (sic)"

(January 27, 2018). Various observations about this heat map generated a significant amount of chatter that quickly escalated between academics and journalists that swiftly brought together the reveal of what Sly of the *Washington Post* reported on January 29, 2018, summarized as follows: "U.S. soldiers are revealing sensitive and dangerous information by jogging."[3] Indeed, the predictability of a running routine or the typology of seldom unexpected activity on U.S. military bases has been highlighted on a global scale. The availability of this data created additional risk that extended well beyond the U.S. military to include Russia reported Sly. A former British officer Nick Waters[4] tweeted: "Patrol routes, isolated patrol bases, lots of stuff that could be turned into actionable intelligence" (January 28, 2018), and through reasonable deduction goes onto interpret further; "So, using this data, you may theoretically be able to individually identify who was, say, doing laps around the perimeter of a possible CIA black site" (January 28, 2018).

Would altering routine in the middle of the desert made a difference, maybe so? However, knowledge of the potential to track data to create a map is certainly not farfetched. GPS data can be used to track individuals anywhere there is a GPS signal unless that signal is turned off. Monitoring of individuals has reached a very high level, potentially compromising sensitive information to the enemy. Disconnecting these monitoring devices may reduce user ease while simultaneously increasing security and decreasing risk.

Typology 7, the *gluttonous optimizer*. This falls into the category of an individual that is motivated by personal gain. Perhaps not to the level of trying to disrupt or place revenge on an individual or organization, but may simply see the organization as a means to an end.

The gluttonous optimizer is likely to display narcissistic tendencies. Dr. Heitler, a practicing therapist, describes that narcissism creates difficult relationships and may be "redefined as a deficit in *bilateral listening*"[5] and instead hearing primarily their own thoughts and brushing aside the thoughts of others except for when an individual wants to impress someone. Then they can demonstrate an aspect of selective listening in order to impress and be recognized.

Typology 8, the *unknowing curator*. Just because someone is the holder of information, or retainer of knowledge, it does not mean that they have a complete understanding of the use of the knowledge

either for its intended purpose or otherwise. Indeed, the curator could be holding information that is the final puzzle piece in a very important puzzle. Release of this knowledge could have cascading impacts and consequences. Also the loss of this information could be equally damaging. The curator could be holding many different parts of a larger puzzle where assembly of all these pieces could be very devastating to an organization's proprietary information.

Typology 9, the *defiant disregarder*. This individual will knowingly violate security practices or training, although their motivations may not be malicious. Working around specific security practices might in fact get the job done, for a boss who is in a rush, or a budget that is not adequately funded. The "yes I can do it" person might fall victim to just too many shortcuts and work-arounds to provide a solution that is less than secure.

Typology 10, the *technically obsolescent*. There needs to be a balance of expertise within an organization where personnel have access to technology refreshes and advances. Investing in the retention of good employees as well as the investment in their training so that they remain current and relevant is very important. It is very easy to underestimate the emphasis in this area. It also means that to remain on the leading edge, experts should attend conferences and panels, as well as read emerging technology literature. When IT communities within an organization are staffed too thinly, then this resource gap or personnel shortage will likely limit opportunity to identify new investments, enhance processes, and reduce risk in information technology security.

Typology 11, the *hygiene-hindered communicator*. Poor receiving and sending hygiene. For example, not ensuring that signature certificates are current and sending to unconfirmed agencies, persons, offices or individuals before verifying the legitimacy of a receiver greatly increases organizational risk. A hygiene hindered communicator may fail to take the time to password protect a document that needs to be protected. Even when there is guidance in place on how to send email, the hygiene hindered will fail to comply by not taking the extra time to verify their work. When snail mail was once the primary means of communications, the majority of correspondence was internally screened for sensitivity before being mailed. Official correspondence was duplicated for hardcopy file retention before the

original was routed for postal delivery. Organizations that put into place selected templates may mitigate their risk of rogue messaging; however, this will not protect misdirected mail. Also ensuring that email can be audited could help detect and identify potential high-risk areas.

Typology 12, the *goofing around player*. The best example of this abuse I can provide is someone who sees an open computer, jumps on, and sends an email to themselves or others that is not appropriate, then returns to their own computer responds back to this fake email for record. However, what this demonstrates is that having an unlocked computer terminal can bring about a large number of risks depending on what is stored or managed by that particular employee. Additionally, jokes of a harassing nature could have a counterproductive intent. Goofing around in an office setting with computer technology is simply a bad idea.

Typology 13, the *blissful user*. If employees—computer users—have a lack of awareness and training, they may not be aware of the risks they are creating for the organization. Privileged users, especially if they are socially engineered, might unbeknownst to them allow perpetrator(s) into the network that might stay and observe for some time to better understand the architecture in order to do damage in the future. Blissful is just unaware and is not seeking out an awareness. It can be risky for contractors to be blissful users as well, as less of their work may be subject to internal monitoring if the contractors, or subcontractors, work off-site.

Typology 14, the *trusted believer*. It is best to be a skeptic when it comes to insider threat. Individuals that take things at face value, or have less-critical thinking ability, might be more prone to social engineering attempts and remain unguarded. It's important to remember that everyone can be tricked or become a social engineering target with even more precision in a spear phishing attempt. Whereas social engineering was once thought to be a trap for fools, perpetrators have created a false sense of security with lures from princes abroad and relatives stuck without money on a vacation gone bad. However, this modus operandi has changed a great deal.

The Federal Bureau of Investigation's Internet Crime Complaint Center (IC3) continues to track the methods used for exploitation; the IC3 recently issued information on the $5 billion dollar scam[6] of

email account compromise that targeted businesses working with foreign suppliers and/or businesses that regularly perform wire transfers. Individuals that routinely perform these transfers have been more narrowly targeted. The entry points have varied depending on the attempt that is made toward the individual. The individuals within an organization that have the authority to conduct financial transfers should have additional training on what trust and verify actually means, given specific social engineering attempts. IC3 has identified other forms of email fraud that include promises of romance, unclaimed lottery winnings, employment, and rental scams, all targeting unknowing money mules. Spoofed high-level executive accounts make this easy, as well as attorney impersonations, fraudulent invoice payments, and data theft. Without established intrusion detection systems, emails that appear official in nature may "leak" in and be incorrectly trusted.

Typology 15, the *rushed surge responder*. During surge operations and/or crisis operations, information assurance may decrease leaving gaps between traditional security, such as the physical, and information technology security. External political factors can push an organization into recovering quickly without taking proper precautions to protect or assure information safeguards. Organizations can have responders on retainer, essentially set up to respond to a variety of crisis scenarios, so that the crisis can be better managed without the added pressure of contracting negotiations. The number of crisis is vast but could include a false claim that has gone viral and could jeopardize company trust. Forward planning for an unanticipated crisis is critical. Having technology that can support surge operations that is current, has been updated and is ready to go and is trackable in inventory is also important. Mobile technology in the case of a lawsuit may need to be recalled. What mobile devices are authorized, and which devices are to be used may need to be part of the preplanned surge strategy response plan.

Typology 16, the *memorable password scribe*. Passwords are numerous and can be difficult to remember, and with increasing length and rules for each password type, individuals are prone to write them down. When passwords are written down they have a greater chance of being lost. Remember that previous Uber example: Loss happens everywhere. Also, passwords are more likely to be repeated for different websites, products, or accessibility. Individuals may even email them back to themselves to remember or take pictures of their

passcodes. Both work and personal passcodes can be similar and could allow for more accessibility via spear phishing. Maintaining safeguards for passwords, changes, and resets are important to security protocol health. Writing down passwords and keeping that book in the right-hand drawer of your desk is a very bad idea.

Typology 17, the *privileged user abuser.* This particular typology is an example of the ultimate MacGyver, a person who is talented and is a privileged user and is able to create a work-around with permissions. Whereas work-arounds can be of benefit to getting the job done, these work-arounds can also create loopholes where a MacGyver might lower security requirements to stream a pornography movie on their computer, or allow computer gaming, or support other personal interests online that are not authorized.

Typology 18, the *unmonitored teleworker.* This category goes beyond the traditional work productivity at home, and instead goes to information security hygiene at home. It is possible that there can be other members or visitors to a household that should not be aware of what an employee is working on. The same level of security protections may or may not be on an employee's home computer, if that computer is used for professional work could potentially expose the organization's network. If an employee brings a work computer home, there is also the increased risk of theft by leaving equipment in unlocked cars, or the risk of being stolen in a home burglary, for example.

Typology 19, the *wide-open wanderer*. This typology has been touched upon as part of some other behaviors. You know this person, I know you do. It's the person that gets up to go to the bathroom or the printer, and doesn't bother to lock their computer, leaving it open for anyone to log into. This can happen at any level of the organization, from entry level to executive level. I've seen this happen even at service organizations where I know there is a lot of personal information stored. Employees should immediately lock their computers when they get up from their desks, no exceptions, no excuses, and period.

Typology 20, the *verbal space cadet*. Oh boy, how I love to hate this particular typology. I have seen this typology increasing in recent years, especially when some mobile devices now require special adaptors for privacy headsets. For whatever reason, many conversations are now being placed on speakerphone in open places with increased volume. Not so smart. Many topics of conversation are not private

anymore because of this verbal space cadet. Ok, so I'm guilty of starting to engage in some conversation when the phone is on speaker, just to raise the awareness of the mobile user. It usually meets with negative looks. Those who use the speaker phone that ignore their physical surroundings and are not aware of what personal information they are disclosing, while the individual on the other line has no idea they are surrounded by strangers in their conversation, places everyone at risk. Information that needs to be protected is especially at risk. In the office environment which is tragically set up with modular cubicle furniture, despite the best acoustic set up, is still at risk. It is best to think about what personnel need to be assigned to offices, what team rooms there are, and what protocols are in place for discussing specific information, a when and where. As an exercise if you spent just one day listening actively in the space around you, throughout the day, you are likely to identify potential threats and even actual cases of increased risk.

Typology 21, the *unguarded doorman*. This typology might be viewed as a not-my-problem character. For example, the credentials should be checked for the technician that is reviewing the cables in your closet or upgrading particular software in-person. There are entry points into networks and if unguarded, intrusions are more likely to occur. If an individual wants to take virtual control of a computer in order to fix a problem, and the unguarded doorman authorizes this work without checking legitimacy and credentials, this poses an increased insider threat.

Typology 22, the *absentminded*. This typology is worth discussion because it is about the absentminded, perhaps the good employee that might seemingly forget, and it's a one-off thing. Stress can lead to forgetfulness as well. It's important to require vigilance of all employees, and if some employees start to deteriorate in this capacity, it might be time to retrain or assess the employee–employer relationship. Medical situations can emerge that might create this type of circumstance. Situations like this do occur in both the private and public sectors. For example, performance may deteriorate because of accidental injury on the one hand, and decrease mental alertness from dementia or another medical condition that might require medication that brings on fatigue or mental fog on the other.

When an unfortunate situation like this emerges, individuals should be assessed for readiness and level of position. A position may

need to be adjusted or shifted into a function with minimal risk and increased supervision. Adherence to Civil Rights law needs to take place so ensuring that action is taken that meets legal requirements is equally important.

Typology 23, the *face-off subordinate*. Who hasn't observed this once in their careers, right? The face-off subordinate is about those employees that do not like their supervisors and do not make life easy for their supervisors who look out for their best interests, essentially outright sabotage or withholding vital information. The subordinate may still be true to their perception of ethics and a love for the organization, but are unable to really view their behavior in terms of not being able to see the forest for the trees. Employees that are toxic to their supervisors are also being toxic to their organizations (and not to say it doesn't work the other way around—it does); however, an employee may not realize how their actions are sabotaging the work environment, larger unit, organization, or bottom line.

Typology 24, the *shortcut alleyway takers*. This one is fairly straightforward; those who want to take the shortcut and be on easy street are at a higher threat level to the organization. The easy way could be the least expensive way for technology reinvestment or easier yet, outsourcing a problem without fully understanding it. Neither of these options are optimal solutions as each poses a higher threat level to the organization. There are many contracting examples of taking the easy street and being stuck with a much larger bill, or even outdated technology when all was said and done.

Typology 25, the *uninstructed newbie*. If an individual is not familiar with a particular system or process, they may place an organization at risk without realizing it. Using technology without the proper training is a risk. At a simpler level, if an employee does not have basic cybersecurity training before being given access to a computer network, risks could be introduced. For example, if specific sites are not blocked, and an employee accesses personal email, there is a higher risk of creating an entry point for malicious activity to the network from the outside.

Typology 26, the *surprising superhuman*. Listening to other people's conversations is a form of active listening. In a modern-day work environment, office space is often at a premium for business expenditure. There are many employees in open-concept systems

that have learned to tune out conversations in order to focus on their work. Whereas there are others that instead pay attention to what their co-workers are saying around them. Even the whispers. Some employees may or may not even try to listen, but the environment is too open and inappropriately transparent. Without making a significant effort, a lot of information can be gleaned simply by entering into an active listening mode. Even staying in a stationary location, the conversations often simply walk by the desk, chair, or position of the employee. An active listener does not need to wander around. These conversations will go to the active listener. Many years ago, as a teen, when I studied at the American Conservatory Theatre in San Francisco, one of the acting exercises was to go into the neighboring park and listen to conversations. After listening to various conversations for about an hour, the goal was to select a conversation and perform an improvisation of what was said. The results were hilarious. The variations of conversations that were acted out and the diversity of information gleaned in one hour were quite interesting. Unfortunately for business, the results of this insider threat are less than funny. Even without electronic listening devices or recordings, the surprising superhuman as an active listener can be an insider threat by knowing or sharing information with others which they should not know about.

Typology 27, the *ivory towered*. Those who are elite, or more specifically, internalize their position status as elite or privileged without full regard to security protocol within an organization are a risk to the organization that should not be ignored. This can be difficult when those who must hold the ivory towered accountable are in a position that is subordinate. Elevating the importance of security and cybersecurity compliance to the ivory towered is a difficult task at best, but it must be done in order to mitigate this risk. I assess that had the strength of the security office been elevated higher in the State Department while Hillary Clinton was the U.S. Secretary of State, she could have gone on to become the 45th U.S. president. Sadly, for Hillary Clinton as a presidential hopeful, and as the Democrats' primary nominee for president, her complete disregard for security hygiene gave the opposition plenty to work with, and public opinion on this topic was decidedly against her. Should the security offices have been monitoring her practices and requiring compliance? I would assess they should

have. To a certain extent her employees showed wicked insider threat behavior. Organizing traditional and cybersecurity at a higher level within an organization can assist those in the ivory tower to better comprehend their exposure to risk to then decide just how much risk they are willing to accept, on behalf of not just the organization or business, but at the potential demise of their own careers.

Typology 28, the *road traveler*. A business traveler or tourist of the world may accidentally introduce threat into their computers and mobile devices which can have a cascading impact into a more expansive network once they have returned to their office or home. Business owners will often provide free public Wi-Fi as a service to their customers and as a means to draw more customers into their stores. These convenient Wi-Fi connections are often referred to as hotspots and their presence has exploded in numbers globally during the last decade. Although this sounds great, the reality is that the privacy and security of these open hotspots is extremely weak and at high risk to compromise. Names of these legitimate hotspots can be modified just slightly, so that unsuspecting road travelers are really entering the internet through a malicious hotspot that will track their every key stroke and mouse click or tap.[7] In contrast, the legitimate hotspots provided by business owners will most likely be lacking of encryption, be more susceptible to software vulnerability and be more open for eves dropping, often called snooping and sniffing or a man-in-the-middle attack where transmissions are basically run through a reader for a third party to review.[8]

Businesses can establish Virtual Private Networks (VPNs) for their employees to mitigate this risk. There are various software companies that can facilitate investment in this technology. Additionally, having technical support available to guide employees through the complex settings of their devices to minimize risk while traveling, is important to maintaining good security protocol. Many popular tourist attractions are also popular virtual pickpocket destinations. From a personal risk standpoint, never visit financial websites that contain your financial information from these Wi-Fi hotspots. Simply browsing for the fun of it or posting updates for social media should be avoided when a secure network is not available. Fun pictures will look just as fun if there is a delay in posting and a Wi-Fi hotspot is the only means for posting.

Typology 29, the *part-present part-timer*. This could easily be a full-time employee who has their mind elsewhere and believes that they are conducting other work of greater importance or value outside of the work environment. A part-time employee who is also relying on more substantive work in another organization could be focusing on a larger payday elsewhere. The back and forth of information between work accounts to private accounts should be avoided—at minimum monitored, as well as potential misuse of proprietary business property.

Entrepreneurs can easily have multiple ventures underway, but ensuring that there is a clear delineation that does not violate policy or ethics, and what the lawyers call a *Chinese wall*, where the conflicts of interest shall not meet. Proprietary information could easily be syphoned away from the primary organization. Greed can factor into this, but also not paying an employee enough money to sustain themselves where they must have two and three part-time jobs or ventures can create a risky situation. Is there an organizational policy on the types of work that can be conducted outside the organization? What type of non-disclosure agreements should be in place? These are questions that should be addressed periodically.

A part-time employee might not invest enough time, or be allotted enough time to understand organizational protocol, procedures, policies that ward off insider threat, or realize they are part of the equation. A part-time employee might more easily dismiss rules because they are not invested as a full-time employee.

Typology 30, the *storyteller, too-busy-to-tell*. There are many lessons learned within an organization, some stories need to be scoped and should be shared. There can be resistance from disclosing what is perceived to be a negative. However, some could be more strategically shared so that the reputation of the organization and individuals is not damaged, but the lesson is ingrained into the culture. Opportunities for success can be created through organizational storytelling. Traditional security and cybersecurity professionals are, for the most part, not used to storytelling within their organization, and unfortunately mistakes are often repeated. Individual accountability is important, but sharing these lessons learned and best practices through creative storytelling can last and shape a culture.

Typology 31, the *not-checked-out employee*. Today's worker is much more portable, and their role may become interchangeable within a particular industry, even within the same organization. If an employee changes roles, divisions, functions, or converts to a contract employee, keen attention should be placed on the individual's privileges and accessibility. If levels of access are no longer required to access particular parts of an organization, either through traditional security or cybersecurity, then these privileges should be adjusted or eliminated. I once observed a former executive of an organization acting in a different role as an entrepreneur in an attempt to obtain proprietary information from a junior member of the organization. While this junior member held up to the pressure, I'm sure that it would only be a matter of time for the former senior executive to get the information sought from someone who perhaps would not hold up as well.

Recently, I was visited by a former employee who had gone through the chameleon change and I wanted to expand on some of the progress I had accomplished as well as recently discovered risks. However, because I knew this person's role had changed, I had to edit my dialog and pay very close attention and choose my words carefully. Without this conscious effort, I could have disclosed information that was not intended for the audience.

Typology 32, the *laissez-faire trainee*. All employees in this modern era need both traditional and cybersecurity training. Without monitoring and accountability of training completion, and testing and evaluation, the benefit of training can be marginal if an employee does not understand fully the reason for this training or does not absorb the materials. Memorization without more complex understanding and pragmatic application leaves a higher risk to an organization. In some law enforcement agencies, a letter with police officer acknowledgment is placed into the officer's record upon completion of training. This places a certain amount of personal accountability on the officer when operating within the line of duty. As cities and states often bear the financial burden of errors in the field, this accountability demonstrates their diligence. If the police officer does not properly follow policy, especially in more extreme cases of disregard, the accountability and liability could shift to individual police officer as an individual and not covered under the scope of the city or state.

While these are 32 typologies presently identified, as time progresses, more behavior typologies may start to emerge. Training that once provided support, including on-the-job training, may gradually disappear as more individuals work in different ways. If organizations use less employees and more contract support, the risks can also change or shift, especially if the work is performed off-site.

It's impossible to predict the future. The projections for combat and for security were quite different at the turn of the millennium than they turned out to be. The global war on terrorism changed the approach to global security and as a result the world changed its strategies and resultant cascading policies, regulations, and enforcement methods. The United States underwent a massive federal reorganization, the largest since WWII, in which the U.S. Department of Homeland Security was formed.[9] Women became accepted in traditionally filled male combat roles. Several of the transportation systems sector modes of transportation had to change their security approaches significantly. The awareness of threats in cybersecurity have gradually increased, but there is a lot more understanding that needs to be gained for protection against these threats. Insider threat is one of these cybersecurity threats, and there has been less awareness of these threats. The 32 typologies of human behavior are an example of very specific actions where action can be taken to mitigate these risks.

Twenty years in the future, what will the workplace look like? Technology applications are increasingly changing. Even today, a seemingly harmless application may carry with it higher levels of risk; for example, the fitness tracker when geographically placed on a global display. Showcasing how workers seemingly born with a smartphone in their right hand are more likely to feel comfortable with using this technology while not likely fully understanding its risks. Technology has become so integrated that the connection aspects, even to technological administrators, have become increasingly abstract.

3.3 Organizational Process Risk Factors

Processes within an organization should be known and clearly understood by those who must operate within them. This seems simple enough, but it is not always the case. As time marches on and employees rotate in and out, and new policies supersede older policies, the

understanding of processes in organizations can become increasingly murky. If training does not take place to explain these changes and follow on measurement and accountability is absent, insider threat will greatly increase. However, not all of these typologies can be mitigated through computer settings. Validation and verification are often needed ranging from audit controls to human verification of information. Policies and procedures should be in place and adapted, based on assessments, in order to mitigate risk. Reaching out to employees on a regular basis will be important to ensure they understand these communications and that a trusted relationship is fostered and developed. For example, if computers have standardized configurations, organizations can change logon screens to reflect areas of cybersecurity emphasis. Employees should feel that their observations will be taken seriously. Monitoring is important. Agreement, or better yet, a reminder of monitoring and compliance with policy can also be placed on these initial login screens.

The larger the organization, the more complex these processes become in practice. Processes can also be informally changed by managers directing changes to these processes, but without the proper updated guidance, risk is introduced to the process which quickly increases misalignment. Insider threat, the virtuous, wicked, vengeful, and malicious, can breed in these types of murky environments.

Various processes should be considered to counter insider threat. The following is not an all-inclusive list, but it's a good start: audit; external communication outside the organization's network; policy health; prioritization of investment as related to risk; training and education; chain of custody; management of behavior risk; communications of risk to senior leadership; and organizational reputation considerations.

Audit systems should be in place in an organization. This includes conducting audits of traditional security as well as cyber-related audits on a random and routine basis. Organizational leadership should ensure that audits are a catalyst for organizational change and should follow through and monitor those changes. Audits can be conducted internally or by external sources. There are many industry standards and benchmarks to measure against, depending on the process or system being audited. Without audit systems in place, unchecked errors will perpetuate the insider threat.

Sending data outside the organization's network becomes an issue when the information either arrives at the wrong destination or the content that is being released outside the organization violates policy, both perpetuate the insider threat. It is more difficult to assess that information has been lost in this way.

As part of the organizational process, policy plays an important role in establishing practice; however, if some of these policies are perceived to be too difficult, and there is not sound enforcement and accountability, work-arounds are bound to quickly emerge.

Investment in the correct organizational prioritization against insider threat is not always determined at budget boards. Competing demands may take priority. Although the workforce is becoming savvier about insider threat, a knowledge gap remains in cybersecurity. Even with the prioritization of increased training and education, the exploitations of and through information technology systems continue to grow. Real-time awareness needs to be pushed on an ongoing basis. Loss of sensitive data resulting from poor chain of custody hygiene greatly increases insider risk.

Without impeding on civil liberties, organizations should know their employees. When personnel behaviors start to change, managers are in a position to refer employees to employee assistance programs, or security managers, depending on the nature of the change. Throughout an individual's career, it should be expected that there are ups and downs in each person's life. However, some of the downs could leave an employee more vulnerable to being socially engineered, or susceptible to bribery, or even engaging in other vengeful or malicious activities against their organization. Both traditional security and cybersecurity risk should be looked at. Incidents of workplace violence and active shooter situations could emerge from behavior that should be red flagged. Even domestic situations could escalate into workplace violence. According to the Census of Fatal Occupational Injuries, nearly 2 million American workers have reported being a victim of workplace violence and statistics show that reported deaths in the U.S. workplace also continue to increase.[10] Workplace homicides increased to 500 deaths in 2016, as did workplace suicide to 291 deaths, the highest figures since 2010 in homicide and the most suicides since reporting began

in 1992.[11] Protective service occupations increased by 68 fatalities to 281 deaths, an increase of 32 percent.

On April 12, 2012, James Wells, a disgruntled U.S. Coast Guard civilian employee shot and killed two co-workers, firing multiple times with a .44-caliber weapon at his place of employment at the U.S. Coast Guard Communications Station Kodiak, Alaska. It was posited that after increased supervisory oversight was placed on him and was denied travel orders to attend a national conference as a subject matter expert on antennas, that he committed this heinous crime. It was a conference that he had typically attended each year and it was deduced that the escalation and planning of killing his colleagues had started as a result of his perceived exclusion.

Although James Wells did not testify during his trial, he maintained his innocence. The evidence convicting him was convincing, including finding matching .44-caliber rounds in his home, establishing access to a weapon (that was never recovered), a flawed cover-up story that pointed to fake telephone voicemails to both victims in the office, and a self-inflicted nail into his tire as a cover story. He was convicted beyond a reasonable doubt of killing retired U.S. Coast Guard Chief Boatswain's Mate Richard Belisle, and U.S. Coast Guard Petty Officer First Class James Hopkins. Wells was sentenced to the maximum of four consecutive life sentences and is presently serving his sentence in a federal prison in Colorado. However, in December of 2017, the 9th Circuit Court of Appeals granted Wells the right to a retrial. A panel of judges ruled that prosecutors repeatedly overstepped their boundaries including using an expert witness in workplace violence; a new trial has been ordered, though Wells remains incarcerated pending trial outcome.[12]

Workplace incidents are not without other consequences, such as reputation and liability. As summarized by Wireless Estimator (which publishes an industry related communications news blog), the spouses of Belisle and Hopkins in 2016 sued the U.S. Coast Guard under the Federal Tort Claims Act asserting that the organization should have known James Wells was disgruntled, dangerous, and out of control citing the evidence of numerous reprimands and disciplinary sanctions.[13] When an incident occurs that causes loss of life, and in this case homicide, it is human resources records that can be looked back upon to determine if there were opportunities to identify the threat

earlier, and to take steps to identify behavioral issues to minimize the potential for a crime of this nature.

Damages created through cyberspace can have consequences that are financially damaging as well as potentially creating a physical threat that could lead to death. A SCADA system, a pathway to be able to communicate a potential increased risk to senior managers and/or executives should be available. If security is a lower priority in an organization and not immediately available in a crisis then security delays can be costly. The perception that risks in traditional and cybersecurity along with perceived potential for insider threat exist within an organization, an organizational reputation can be easily damaged. If good employees perceive that an organization is no longer a good option, they may look to transport their careers elsewhere.

3.4 Physical Environmental Risk Factors

Traditional security, including physical security needs to be considered in terms of the role it plays in cybersecurity and insider threat. Illusions of privacy and security are rampant. Newer buildings with office cubicle spacing can have even less privacy and security with potentially alarming open access. Physical barriers may not always be in place properly. For example, if a locked cabinet is in place to secure a server, but the doors remain open, then physical access to cyber connections is accessible with ease. These connections need to be guarded. Wiring in buildings runs through walls and can be contained in specific closets designed to house connections. However, these spaces could double as storage, or janitorial supply closets, leaving them open and exposed shows a potential threat. These types of spaces may not be under surveillance as are larger areas such as main passageways and security checkpoints. Additionally, if storage containers for materials, proprietary or otherwise, are subpar, this will also expose an organization to vulnerability.

Some observant employees pay attention and often know where the gaps in security are and may be willing to share these vulnerabilities if asked directly. For example, I frequently see workers in various employment settings such as fast-food chains and other retail settings huddling in the blind spots of security cameras to access their phones and reach their social media network. Probably because I tend to look

at security problems, I may ask workers questions they really shouldn't answer—and surprisingly, employees will share with me the vulnerabilities of their security systems.

Unprotected equipment, especially computer and mobile technology that can get access or allow access into a network is a risk. Open terminals allow a passersby to download files or upload any number of malicious malware. Discussing progress of programs or projects within the close proximity to those who do not need to know should be minimized. The placement of screen views should be considered. In a workplace that feels comfortable, it can become easy to slide into complacency, to not pay attention to who is walking by, who is keeping track, or not keeping track of computer equipment. It's easy to forget that distance in cybersecurity does not matter. An internet connection thousands of miles away is closer than walking across a room to open the door. Once information is sent, it may route through a server that is repeatable throughout the world, in mere fractions of a second.

3.5 Architectural IT System Wellness Risk Factors

Discovery of wellness risk factors of an information technology system may require both a remote and physical validation and audit. Even if policies and rules are in place, the common organizational practice needs to be clearly identified. If you are in an organization with a distributed workforce, with locations in multiple areas, it is likely that consistency of your infrastructure could vary. The openness of systems needs to be examined, as well as validating the people that should have access to these systems, and the identification of those individuals that really should not have access. Many organizations will lock down computer ports when terminals are unplugged, reactivating them when someone is authorized to connect. This takes additional time but can minimize the risk of unauthorized access into a network. As discussed, simply connecting a mobile device into a computer to charge the battery creates a connection that can allow for the transfer of data and a doorway will open. Ports should be labeled, or excess ports be blocked as a precautionary measure.

Individuals that are given increased privileges or access should be scaled and actively managed. Vetting and regular revalidation of the need should occur. For example, if someone only needs high-level

access once a year to submit a report, they should probably not retain the access all year long. It is just not good security practice to do so. This access should be provided for a very short period to submit the report and then be closed. In some cases individuals are given proxies to do work, and these proxies should be used on a case-by-case basis and not be retained. Individuals that transfer positions within an organization could be a risk for the organization if they continue to carry the accessibility controls and are tempted to either use them or are complacent allowing someone else to access. I witnessed the spouse of a worker a few years ago accessing an application from home in order to change some personnel data, in a web-based system—a system that she should not have been authorized to use or navigate in. She navigated the system with ease, and this surprised me.

While the data changes were ultimately authorized, the spouse did not have organizational permission to be in the system; the credentials of the worker could have been used to perform other changes and view other information had they gone to different viewing screens. Without more advanced biometrics, this might have been difficult to detect, since the two factor authentication credentials were present in order to access the system.

When a user has root access, a lot of damage can be done because the user has elevated privileges that can delete and move around infrastructure. If a malicious threat is in your network, deleting may not be the aim, and instead the aim may be to gather as much information as possible. Terrorists and malicious users are not always in the game to destroy, but to collect. With root access, once inside the network, they can easily hide for months because they would have used those credentials to build authorized and legitimate places to hide. This can be very difficult to detect and explains why some companies release breaches well past their expiration dates, or the breaches become increasingly vast than initial reports. Computer forensics can be difficult to achieve.

Technology that is seamless can also cause problems. The future workforce will likely rely more on technology and will have increasingly technology habits (good or bad) that they bring with them from home and college into the workplace. These practices not necessarily cybersecure, which place organizations at greater risk. Quite frankly these changes can be difficult to keep up with. Their patterns and

weaknesses may be more predictable, as the latest technology seems to really be attractive to trying. Technology applications are continually changing; a seemingly harmless application may carry with it higher levels of risk. Workers who have been born with technology in their hands are more likely to feel comfortable with using this technology, and not likely to fully understand the risks of it because it is seamless. Technology has become so integrated that the connection aspects, even to technical administrative workers has become increasingly abstract.

3.6 Aggravated Risk Mitigation Approach

It is important to not make the situation worse, and therefore a plan is needed. Without a plan to respond to an elevated threat against a traditional security or cybersecurity incident, you could be making your situation even worse. You can soon find yourself responding to alarms where you may not be able to determine if you have a real insider threat or if you are dealing with nuisance alarms. Meaning you can't see the forest for the trees, especially if you do not know what your network looks like, or who is connected to your network at any given time. These shorter term and often false alarms can be distractions from a bigger issue that needs a larger organizational change action. For example, if you have an intrusion, it is possible that the intrusion is not a critical one and taking swift action could interfere with critical operations.

However, there may be other approaches to determine the level of the threat, and to determine who is in the system, how they got in the system, and why and/or how the intended breach originally occurred. For example, assessing if the attack was deliberate and targeted, or random and/or serious, or a nuisance, all should be a priority. The scope of the compromise will need to be determined as well. It might be possible to isolate the breach without alerting the intruder to prevent further compromise before disconnecting. It might be possible to place some false information before closing the door. Some very good technical experts are needed for this counter-operations practice.

Marketing and damage control are very important. Mapping accountability chains will also be important. If an investment was not taken to save money, but the intrusion cost the organization more

than the original mitigation measure, this should be known for future investments. Making a problem worse through a haphazard response approach is not ideal, and this approach should be avoided. If the network and technology connections are monitored, and a threat is identified, understanding exactly how the network is mapped will be very important to isolating a threat. If a large intrusion has been identified, the gut reaction might be to launch a full response. However, a quick pause to come up with a course of action is recommended; or a preestablished course of action that makes sense could be implemented; in either case, the response should not cause larger consequences. It can be both a deductive and inductive process depending on the multitude of issues. Combining both traditional and cybersecurity assessments will be needed. Because technology is not always in the clouds, and there are physical connections, these connections need to be audited and/or verified. Understanding if a connection is isolated and whether there will be secondary impacts to an even more critical system will be important to know. There may be multiple solutions, involving various mitigation techniques, and the more that is known the better so that risk can be reduced with the best approach. If your network is not mapped, that is a very good place to start. From an information technology perspective, it will be very important to have mapped your network so that you can isolate potential risk areas if needed. A physical inventory along with updating the latest anti-virus protection is important. Also, if your employees connect straight to the internet and are using standalone computers, you should know what networks your employees connect to. So, how well is the information that they are storing for you or the proprietary information of your/their clients, protected?

In the physical world if an environment is not safe, it can often be assessed through a boots-on-the-ground approach. Assessing everyone, from the lowest level employees to the most senior, including contractors and other personnel, with connections. The insider threat response needs to employ dual approaches. If there are domains on the network that should be removed, or old software that pose a risk, or out-of-date operating systems, these should have been identified and removed in advance. However, responses to situations can in hindsight bring these types of mitigation strategies to light—typically and unfortunately, later than they should have been.

If the operator of a computer terminal, through their risky behavior, created the open doorway and was indeed part of this insider threat, then on-the-spot remediation training will be needed with the individual, or other predetermined course of action. Accountability remains important. Taking this lesson learned and knowledge to determine if the same issues are occurring more broadly with other computer users on multiple terminals will be as equally important.

3.7 Unknown Factors

Recognizing that we can't know everything is important because factoring in unknown risk is critical to recovery. Individuals and their organizations may be trying to make the right decisions but with limited visibility. The unknown is always a risk, because you don't always know what you don't know. Even after an incident, gathering a chain of events leading to a casualty can be a challenge. Supply chain sources can be questionable even if an organization has a trusted source, the threat of a product may change relatively quickly, leaving a security gap. To have a high level of confidence, an entire supply chain must be verifiable; from the manufacturer to transportation to the validated vendor that sells the part to the installer.

I took a systems thinking course with a very dynamic group of federal aviation administrators at the Department of Transportation's Center of Excellence several years ago. Collectively, we explored taking a "systems thinking approach" in order to reduce some of the blinders or ideological perceptions. What we discovered is that there are a lot of perception risk areas. This course I think was a catalyst for me in researching the concepts of thinking critically.

3.8 Practical Mental Models Continued

As discussed, mental models are beneficial for conceptualization of complex topics. A more comprehensive understanding of insider threat should begin to emerge. I continue to unfold the 10 mental models that are all aspects of a grounded theory of insider cyber threat. Building on the prior two mental models—*Aspect 1: At the Crossroads* and *Aspect 2: The Virtuous and the Wicked*—I will now introduce two more. The chapter will continue to build your understanding of this

definition and conceptualization with very practical application. The creation of a new understanding of insider cybersecurity threat to an organization is being revealed to you.

3.8.1 Mental Model (Aspect 3)—Risky Human Behavior as Typologies

The next mental model I call *Aspect 3*, which is *a grounded theory aspect of risk factors and unintended insider threat; 32 human behavioral risk typologies*. To capture the human risk behavior emerging from the data, typologies of human risk were grouped to illustrate groupings of organizational human behavior. Risky human behavioral actions were captured, written into more general categories, and organized into typologies that actually create mental models. These mental models contribute to an organizational understanding of unintended insider threat. These specific 32 typologies that are factors of insider threat may also stand alone as mental models as their descriptive name indicates; however, for discussion purposes, I have captured them categorically.

Table 3.1, specifically Core Category *A,* includes a brief description for each typology as a mental model. This may provide insight and potential opportunity to determine applicability for an organization. Some of the behavior's intent may be inherently clear, while other reasons will need to be later explored by an organization in greater depth so that organizational training or policy changes are addressed in a specific organizational environment. Ultimately, the *Mental Model (Aspect 3)—Risky Human Behavior* may be drastically mitigated with a goal of vulnerability elimination. In essence, through the creation of mental models, what I have captured is storytelling, a risky behavior folded into human caricatured typologies that espouse the reality of insider threat created by both human action and inaction.

This is a great place to begin your organizational assessment and determine your risk by comparing the table against your existing approach and identifying any gaps or deltas. There is opportunity in identifying potential high-risk areas and reducing this risk through systematic remediation. Risk can be reduced by minimizing the number of risk-behavior typologies in an organization. Awareness of these behaviors is a foundational start for an organization to examine, along with creating opportunities to seek ways to shift behaviors

to mitigate risk. Every employee, including those not typically considered an employee, can create a vulnerability that introduces this threat and elevates risk. The virtuous, wicked, vengeful, and malicious employees all create risk, of varying consequential magnitude, which may not be initially known. The conceptual understanding of well-intended behavior that can also lead to an accident or an unintended consequence is also necessary. This shift of understanding may lift a potential ideological veil for many organizations.

3.8.2 Mental Model (Aspect 4)—The Enforcer and the Responder

The next mental model, I call *Aspect 4*, is *a grounded theory aspect of an unintended insider threat phenomenon of the enforcer and unchecked risk mitigation countermeasure response*. Risk mitigation efforts were shown to be risky, as captured in typology 15, the rushed surge responder, when actions were perceived as inappropriate or untimely without needed collaboration. Cyber threat often requires the support of traditional security, including physical security, especially when dealing with a human entity. Simply shutting down a person's access may not mitigate additional risk that could be created through a physical presence. However, there was also an identified risk of moving too slowly, and that a balanced and coordinated approach was needed to respond to threats. Not to just respond to either the cyber aspect or the physical aspect, but to also be able to assess the threat and manage the response in a more comprehensive and collaborative way between information technology and physical security.

3.8.3 In Review

This chapter has provided a lot of valuable information on insider threat: assessment and mitigation of risks. It answers the question, what are the factors of insider threats to organization, in a readable and pragmatic way. Much of the information presented provided clues into the unintended insider threat, specifically the virtuous and the wicked, and at times the vengeful and malicious insider. As discussed, key organizational themes emerged on human behavior, organizational process, and the physical environment. Notably, action that was

well intended to assist and mitigate risk could also cause additional damage, increase risk, and cause cascading impacts of unintended incidents. A well-coordinated response is necessary for an organization to manage this risk. To do so, additional communications and strategic risk communications between hierarchical levels are required. Unknowns, or unshared known information, were emergent themes and contributing factors to insider threat. Primary examples of factors that contribute to unintended insider threat were presented:

- Human behavioral risk factors as typologies
- Organizational process risk factors
- Physical environmental risk factors
- Architectural IT system wellness risk factors
- Aggravated risk mitigation approach
- Unknown factors

As you may have now identified, with a greater depth, human behavior is a significant contributing factor for increasing the risk within an organization. Insider threat created through risky human behavior impacts cybersecurity and extends well into the physical domain. Along with humans, the organizational processes that humans have created, or failed to create, in conjunction with the physical environmental risk factors are of a concern. How well a broader architectural computer system is initially designed or subsequently modified is a risk factor that also includes particular assigned users, the privileged, the root based, and those with role-based access. More attention should be given to this area. The need to collaborate is only getting stronger; even when employees try to do the right thing, such as identify a threat and take the threat off-line, they can still create a situation even worse than the initial threat.

Unknowns are often not just unknowns in the general knowledge sense, but instead can be known to some part of an organization and to a lesser degree in other parts of an organization. This demonstrates a greater need for retaining knowledge, sharing information, and conducting strategic risk communications surrounding the emergent themes presented for insider threat; assessment and mitigation of risks.

Appendix: Best Practice—Practical Knowledge and Practitioner Application

Practical knowledge

- Review the human behavioral risk factors as typologies, there are 32 identified. Ensure you recognize the increased risks that are created through the risky behavior of the unintended insider threat of the virtuous, wicked, and sometimes vengeful.
- Note distinctions between vengeful, and malicious. At times vengeful consequences are unintended, and not recognized, at other times they are intended. Malicious behavior is destructive and harmful but can also have unintended secondary or tertiary impacts. However, the malicious insider typically falls into the intended insider threat category.
- Recognize that in addition to risky employee behavior in organizations there are other risk factors, specifically: organizational process, physical environment, architectural IT system wellness, aggravated risk mitigation, and the unknown factors.
- Reflect on how the use of mental models may be effectively used to both understand the problem of insider threat and communicate the risk as a tool of mitigation.
- Understand how organization leads for countering insider threat needs should be assigned from multidisciplinary perspectives and work together in collaboration under a comprehensive insider threat program that focuses on the entire threat and risk landscape and not just the predominately external malicious insider threat.

Practitioner application

- Specifically review each behavior typology, identify a current countermeasure in place, or lack thereof, to identify your gaps and deltas.
- Review the factors of insider threats against your identified insider cyber threats to organization Core Category checklist (Chapter 2). Include comparative analysis with organizational

leads, co-leads. Assess and identify any organizational gaps, including missing controls and accountability mechanisms.

- In your organization, assess any recently discovered insider threat events. Assess if the knowledge was unknown, or known and simply not shared. Determine if this communication pathway has been corrected.
- Identify how much of your organizational technology falls into standard configurations and is routinely scanned and patched; identify how to increase the security of the workforce through upgraded operating systems, and standard configurations.
- Determine who are your network enforcers and responders, and if they are available 24/7.
- Review your organization's traditional and cybersecurity training programs and plan to enhance, including adapting stronger personnel accountability measures, at all levels, against insider threat.
- Be prepared to ensure that the people accepting the risk for the organization have the authority to accept this risk, and understand the risk they are accepting.

Endnotes

1. Uber Newsroom. The 2018 Uber Lost & Found Index. (2018). Retrieved from https://www.uber.com/newsroom/2018-uber-lost-found-index/.
2. Nathan Ruser Tweets. Strava Releases Their Global Heatmap. January 27, 2018. Retrieved from https://twitter.com/Nrg8000/status/957318498102865920.
3. U.S. Soldiers Are Revealing Sensitive and Dangerous Information by Jogging by Liz Sly. January 28, 2018. *Washington Post*. Retrieved from www.washingtonpost.com.
4. Nick Waters Tweets. Big OPSEC and PERSEC Fail. January 27, 2018. Retrieved from https://twitter.com/N_Waters89/status/957323495226167296.
5. Narcissism: A redefinition and case study by Dr. Susan Heitler. 2014. Therapy Help. Retrieved from https://www.therapyhelp.com/narcissism-a-redefinition-and-case-study/.
6. Federal Bureau of Investigations. Public Service Announcement (I-050417-PSA). Business email compromise, email account compromise, the 5 billion dollar scam. May 24, 2017. Retrieved from https://www.ic3.gov/media/2017/170504.aspx.

7. What is public Wi-Fi by Norton™ by Symantec Corporation, 2018. Retrieved from https://us.norton.com/internetsecurity-privacy-risks-of-public-wi-fi.html.
8. Ibid.
9. Public Law 107-296. To establish the Department of Homeland Security, and for other purposes. The Homeland Security Act of 2002. Retrieved from https://www.gpo.gov/fdsys/pkg/PLAW-107publ296/html/PLAW-107publ296.htm.
10. U.S. Department of Labor, Occupational Safety and Health Administration. (2017). Workplace Violence. Retrieved from https://www.osha.gov/SLTC/workplaceviolence/.
11. Bureau of Labor Statistics. (2018). Census of Fatal Occupational Injuries Summary, 2016. Retrieved from https://www.bls.gov/news.release/cfoi.nr0.htm.
12. Appeals court orders new trial in a double murder at Kodiak Coast Guard station by Lisa Demer. December 19, 2017. *Anchorage Daily News*. Retrieved from https://www.adn.com/alaska-news/crime-courts/2017/12/19/appeals-court-orders-new-trial-in-double-murder-at-kodiak-coast-guard-station/.
13. Wives of murdered tower techs sue the Coast Guard for wrongful death by Featured News. August 16, 2016. *Wireless Estimator*. Retrieved from http://wirelessestimator.com/articles/2016/wives-of-murdered-tower-techs-sue-the-coast-guard-for-wrongful-death/.

CHAPTER

3

MANAGING VULNERABILITIES IN THE CLOUD

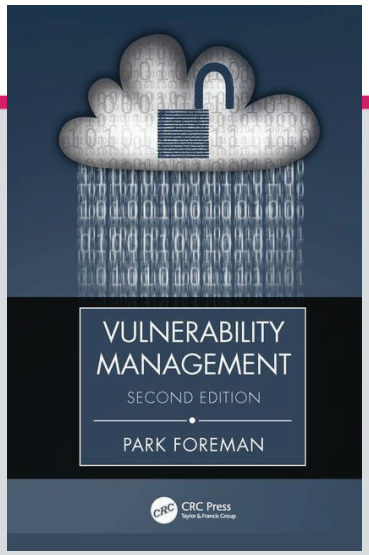

This chapter is excerpted from

Vulnerability Management

by Park Foreman

Learn more

Chapter 10

Managing Vulnerabilities in the Cloud

10.1 Vulnerability Management in the Cloud

This chapter provides insight into vulnerability management where cloud computing resources are concerned. It is intended to supplement the lessons of previous chapters by addressing the issues of identifying and managing vulnerabilities in a cloud world. The basic principles of vulnerability management are unchanged by cloud services. What will change is how the principles are applied.

Until now, I have described vulnerability management from technical and process points of view. These dimensions remain consistent and relevant for those who own or control infrastructure and software. In the cloud computing world, however, ownership and control can become more complicated. Ownership is easily established since contracts of sale stipulate the ownership and scope control of what was purchased. One may purchase software but not control the underlying code components. This is an important point for vulnerability management since the software may need a patch that is controlled and issued by the seller. I have addressed that issue throughout this book mostly with attention to vulnerabilities in a restricted physical environment. With cloud, the customer becomes a renter of various resources. The relevance of this and how vulnerabilities are managed will become clear later. First, it is important to understand the reasons for cloud, which will help the reader understand why vulnerabilities must be managed differently.

Consumers of IT services today have sought to decrease capital expenditures and reduce the potential waste and burden of underutilized assets. In other words, they wish to obtain the required resources on demand and adapt quickly to changes

in the enterprise. The traditional in-house IT model requires one to maintain control of excess IT capacity in anticipation of unpredictable workloads. These IT systems require maintenance and become obsolete over time. So, when the organization requires a new technology that is not supported by the current, underutilized installation, there is potential waste.

Alternatively, there are considerable benefits to an agile technology environment that can respond to business needs in days, hours, and even minutes rather than weeks and months. These capabilities can yield competitive and cost advantages.

To achieve these benefits, cloud service providers now manage computing resources for numerous customers at lower cost and with faster provisioning. When the IT manager needs to deploy a new finance application requiring compute and storage, a corporate credit card and 10 minutes online will have the infrastructure ready. In other cases, the software itself will be available in a production-ready environment immediately and only requires a subscription and the upload of any data and configuration information. Note that there is far less control of the operating environment.

Layers of abstraction are an important effect in the effort to maximize the utilization of resources and lower costs. This has resulted in development of software tools such as "OpenStack" to manage and monitor services sold to customers rather than selling hardware and software. Flexibility and speed have dramatically changed the course of how IT is designed, built, and maintained. Necessarily, the ability to manage vulnerabilities has become more important and more challenging to a larger audience.

When considering vulnerabilities in a virtual world, the previously mentioned "abstraction" can be a pervasive factor in the risk environment. Abstracting technologies can reduce risk and keep costs low, or they can conceal risks since they are no longer visible. Consider the technologies that are virtualized versions of what were once considered to be an immovable foundation. The Host Computer has become a virtual machine. The physical switching and routing in a data center are now a virtual network. Traditional VPNs (virtual private networks) are ubiquitously integrated into the former. Even wide area network (WAN) switching and routing has evolved into software-defined networks.

In computing, most IT professionals have become accustomed to virtual machines where one does not necessarily know what physical CPUs, memories, or disks are in use for a given system, application, or network component. Virtualization has been further diversified by containerization of specific tasks rather than committing even one virtual machine to a single function. Multiple instances of a container can exist in multiple cloud locations, and load balancers distribute or load-balance application requests among these containers. In either case, the results are returned to the requesting application component. Understanding the state of

these containers sometimes sounds like trying to determine which birds in a flock are male or female as they fly around.

Additional blurring of lines occurs when a combination of virtual machines, containers, and software defined networks (local and wide-area) are combined across multiple data centers in diverse and changing geographies. To further understand the challenge of identifying, prioritizing, and remediating vulnerabilities, not even the application itself is all that clear. Many applications share some of the same microservices spread across myriad systems and communicate through a changing, adaptive network. Visibility into this jumbled, shifting virtual world can seem impossible.

So, one does not control most of the resources in use and cannot even consider scanning and patching. Some things can be patched and others cannot. There are two key activities to successfully manage vulnerabilities in this world. First, clearly define the dimensions that are your responsibility and what part belongs to the cloud service provider. Then, implement the appropriate controls for both. *Never lose sight of the fact that you are managing risk more broadly and not just software vulnerabilities.*

So, what exactly is a vulnerability that is specific to the cloud? The "cloud" quickly becomes less mysterious when one realizes that it is just infrastructure, software, and services of varying types that you might find in a traditional, in-house IT environment. But, there is a difference. In order for cloud services to be scalable to many tenants, there has to be orchestration software, shared equipment, virtualized networks, and some specialty software to make it unique and competitive. It is these latter components that become intrinsic to cloud vulnerabilities. To step beyond the straightforward view of scanning and remediating in a traditional IT environment, look for the vulnerabilities that are unique to the definition of cloud. From NIST SP800-145[1], the characteristics of cloud are broadly defined. Here is the short list with some clarification:

NIST Description	*Elaboration*
On-demand self-service	Order through a web page or in-house specialized software.
Broad network access	Internet and/or internal virtual network.
Resource pooling	Share hardware, software, and network platform with other buyers of the services.
Rapid elasticity	Add and remove any resources or technologies you wish.
Measured service	Pay by the year, month, day, hour, or even minute.

10.2 Dimensions and Challenges of Cloud Service Vulnerabilities

In previous chapters, I have described the process of identifying, assessing, prioritizing, and remediating vulnerabilities. If one recognizes that at some point in the dissection of the virtual world there still are hardware, software, and people, then the general process has not changed. Only the participants, their roles, and methods have changed. There may be an operating system patch management role, but it is not in your organization. The firewall engineers responsible for protecting an application may work for an outside service provider. Physical facilities are protected and monitored by several other parties spread around the world and governed in different jurisdictions by different sets of standards and regulations. The techniques for remediating vulnerabilities will vary.

This model is not new. The world has depended for millennia on a complex supply chain and numerous providers of materials and services combined by multiple layers of third parties resulting in an end product that is sold by yet another party. The reader can refresh their perspective by watching the economics lesson in the 6.5-minute video "I Pencil"[2] that can be found with a simple web search. However, there are specific aspects of how cloud operates at a high level that bear elucidation (see Figure 10.1).

Note that there are many parties working in concert, but only one party is held directly accountable for security, the Application (Software) as a Service Provider. The physical data center servers and network are all managed by another party. Yet, the contract between the application provider and the customer at the top of this chain can only hold one party directly accountable for security. Networks, physical security, audits, and so on are two levels removed. This loss of governance and control makes it necessary to break this supply chain down into smaller parts to examine some of the options for identifying vulnerabilities or holding other parties accountable, albeit indirectly.

Generally following the Open Systems Interconnection (OSI) model, a cloud service is provided at one or more of these levels: Physical, Data Link, Network, Transport, Session, Presentation, and Application. However, they are not necessarily wrapped up in a complete package. There are myriad relationships and intervening parties who directly or indirectly provide these layers. All of them can introduce vulnerabilities, and one or more of them have responsibility for detection and remediation.

Naturally, this state of affairs has created some misgivings about cloud services since visibility is limited. And while it is important to manage risk, it is equally important to seize the opportunity found in the competitive advantages of cloud where productive results are obtained faster and more cost-effectively. In fact, the economies of scale are immense. Businesses can adapt more quickly to changes in customer and regulatory demands because there are other experts at work in the supply chain.

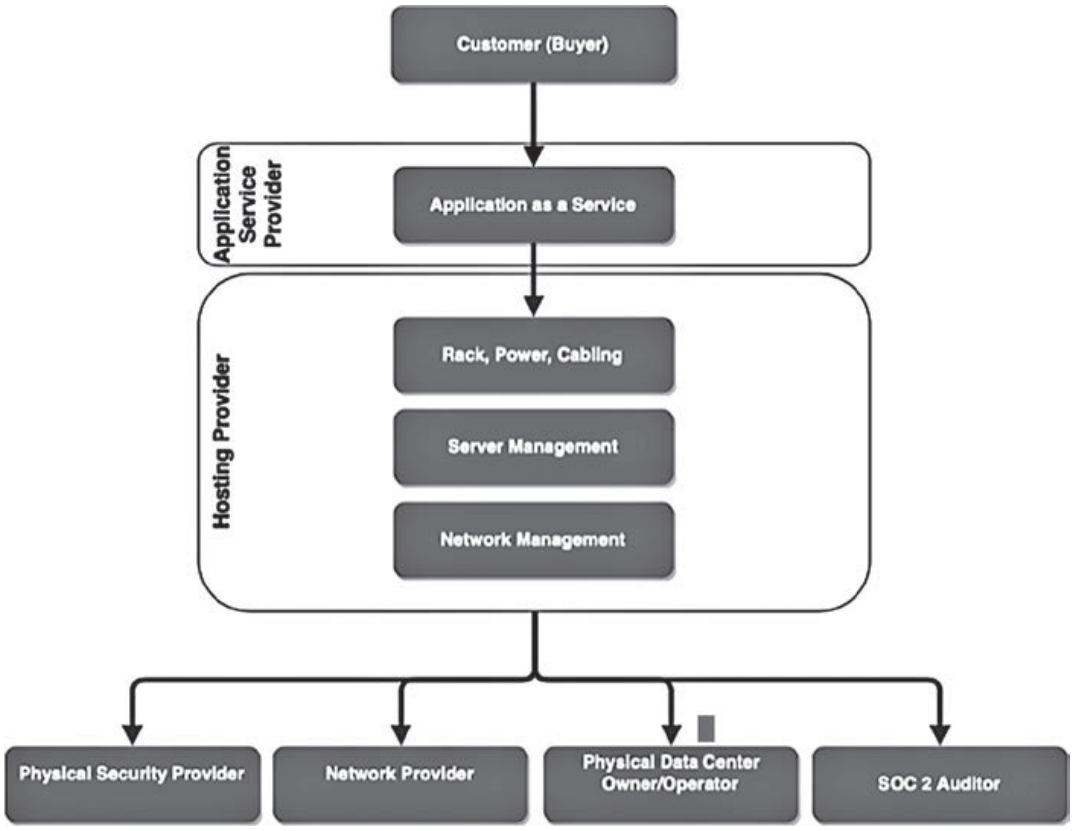

Figure 10.1 Service provider tree.

Furthermore, the buyers of cloud services only need spend as much as is necessary to perform the task at the moment. In the traditional do-it-yourself model, there is a costly amount of planning and build-out of resources in anticipation of a project. This sometimes results in the IT manager over-purchasing to be prepared to rapidly respond to unplanned business demands. The cloud model anticipates this across numerous customers and lessens the "slack" in budgets or plans.

Cloud services are also, in some respects, an equalizer between David and Goliath. The same, cost-efficient services are available to parties small and large. At a very large scale, it may be possible to negotiate some marginal advantage, but only to the extent that it does not become cheaper for Goliath to build his own data center, networks, applications, and software licenses.

All such cloud services can be grouped into the categories: Data Center, Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Each of these has their own risk and vulnerability management challenges that require unique approaches. In the next few sections, we will discuss the broad methodology and some specifics about how to manage risk in each of these areas. With myriad options available, it would be wise to determine the strategic direction of IT and develop a consistent set of processes across all services in order to provide a global governance view.

10.2.1 Risk Management Methodology

Thus far, I have focused on a narrow segment of risk management that deals with technical and procedural aspects of assessing vulnerabilities. For the vulnerability management journey through the cloud, one has to broaden their view beyond scanning, prioritizing, and remediating. Before cloud, operating system, infrastructure, software configuration, and administrative procedures were enforced and monitored for compliance. As explained earlier, in the cloud world, many of these components previously governed by the consuming organization are now partially or fully invisible usually because third parties have taken this responsibility. Early adopters of cloud often assume that all this monitoring and governance is done properly and feel it is just one less thing to worry about.

The latter case is a tactical error for any organization. Lack of visibility creates increased risk exposure or, at the very least, provides no assurance. We will examine these challenges in the various cloud service models and then identify some ways to manage the associated risk. No service provider, organization, or IT strategy combination is the same. The risk needs are as varied as infinite permutations of business operating models. I will describe a variety of tactics in this section. For example, tools such as a SOC 2 report, which is created evidence of the effectiveness of controls in a service organization, can be obtained from any supplier who has chosen to use outside auditors. This report provides some assurances of the confidentiality and integrity of the systems supporting security in the services supplied. But a few tactics are not a panacea for a comprehensive risk management program.

A very useful approach to planning the risk management program for both internal and cloud provider services is to create a matrix of provider security control types and determine the appropriate method of validation based on the risk level. If consistency and balanced decision-making are to be observed, internal activities of your own organization should be no exception and also should be considered service providers.

Begin with a matrix of control types and risk categories for a particular provider. Figure 10.2 shows one possible matrix.

The simple idea is to determine the types of security controls desired from a cloud service where your own administrative security controls can be used for validation. Those controls may vary by:

- the type or classification of data being processed or stored
- business criticality of service provided
- financial impact of a loss event
- legal or regulatory requirements

It would not be unusual to have a different version of this matrix for each category of business data or function. Looking back at Figure 10.1, it is evident that an application (software) as a service, or SaaS, is purchased directly by the customer. However, there are several underlying elements of the service to take into consideration with

Service Control Requirement	Control Validation Type				
	External Audit	Direct Audit	Certification	Contractual	Technical
Identity and Access Mgmt	N	O	N	N	R
Regulatory Compliance	R	O	R	O	N
Data Provenance	O	R	N	R	N
Data Segregation	N	N	N	R	R
Data Recovery	R	O	N	R	R
Monitoring and Reporting	N	N	N	R	R
Business Continuity	O	O	N	R	R
Patch Management	N	O	S	O	N
Security Monitoring	Y	N	R	O	S

R-Required: A control is required for this type of provider.
O-Optional: If required controls are met, this optional control can be used in selecting the provider.
N-No required control.
S-Specific control type is required: A control such as certification or accreditation from an external authority (e.g., Microsoft's "MCSD: Azure Solutions Architect" or Cisco's "CCIE Data Center"). Other certifications such as PCI DSS (Payment Card Industry Data Security Standard) or HIPAA (Health Insurance Portability and Accountability Act) may be applicable depending on the application.

Figure 10.2 Control-type requirements matrix.

the understanding that these too can introduce critical vulnerabilities. For example, the application provider gets network services from an outside provider, yet that network may be the primary transport mechanism for your critical data.

Due to the variety of layers in this stack of services, there are a variety of approaches to managing risk. If the application in question handles confidential data, then it would be prudent to have a matrix governing the secure access to any systems holding or transporting the data. In this case, vulnerabilities may be identified through audits or reviews of security controls prior to contracting. Risks identified can then be addressed through remedial activities such as:

- Encryption of the data at rest and in transit
- Contractual obligations for the service provider to prevent, detect, and report vulnerabilities or incidents
- A contractual right of the buyer to (one or more of):
  - audit the security procedures provided or governed by the contracting service provider
  - receive a security report of controls quarterly or annually
  - confirm the security or technology certification of the provider and their staff

Note that subsequent checks should be performed regularly to ensure continued compliance with stipulated risk controls. One might consider this the administrative equivalent to updating vulnerability scan checks and performing another scan.

But these are contractual and technical controls. There is also the option of using administrative controls such as contracts, direct audits, or verification of annual certifications. Referring back to our grid in Figure 10.2, if we are expecting the provider to have an identity and access management service support of the application, then we could use technical controls such as logging or LDAP (Lightweight Directory Access Protocol) integration to assess some of our own overarching control of that service element. Conversely, where there is a requirement for patch management service, one can apply a specific control related to the service used such as certification in the software or patch management practices. It may also be a simple matter to use the administrative control of having the supplier provide a quarterly report of patching activity or a notification of scheduled patches.

The European Network and Information Security Agency (ENISA) issued a paper in 2009 called "Benefits, risks and recommendations for information security." This paper provides an excellent list of security vulnerabilities to be found in cloud services. In general, the list holds true for cloud and on-premises technology services alike. The report relies heavily on Gartner™ reporting, a Cloud Security Alliance report, and the opinions of several contributors. The overall assumptions are that cloud service providers have common, observable vulnerabilities. Note that these observations are subject to change and variation across a now massive marketplace for cloud services. So, the report is clear in that it provides some broad advice: "The scenario was NOT intended to be completely realistic for any single cloud client or provider but all elements of the scenario are likely to occur in many organisations in the near future.[3]"

One very important finding is the potential and possibly stealthy risk of loss of governance. The loss of visibility and control at each layer of the OSI model can governance and oversight. As a result, it becomes essential to understand the vulnerabilities associated with that loss and formulate an appropriate management plan. As the reader will see in the sections that follow, there are some common mitigating controls intended to address the risk of the "unseen" in so far as there is sufficient assurance of a secure computing environment.

10.2.2 The Data Center

Physical data centers are the lowest level of service that has major capital and operational challenges in the hierarchy of services. The data center provider typically has physical facilities, multiple circuit ingress points (often referred to as DMARCs or demarcation points), video surveillance, guards at a desk to escort visitors, cages for each customer, backup power systems, and sometimes racks in shared areas.

When assessing risk for a data center provider, there are several tools available depending on the services and flexibility of the provider. One way to identify the right service provider is to decide what significant physical risks there may be with any provider and what will be the cost-effective approach to mitigate them. For example, if the data in a company's application are subject to a standard concerning the protection of cardholder data (i.e., in-scope for payment card industry [PCI] compliance), then before making a contract with the service provider, determine if PCI requirements can be met. Can the data center manager supply access logs, video camera footage, and any other PCI requirements relevant to a physical data center holding cardholder data?

If the answer is no, either you will have to find another, more suitable provider or create a control (or mitigating control) to remain compliant. For example, one could place surveillance cameras in the rented cage in the data center and maintain an access log. Alternatively, it may be more practical to use strong data encryption of data in transit and in storage. It would become increasingly difficult for an auditor to argue that there is a physical risk to cardholder data.

If none of these options are available, another approach may be to negotiate a contractual agreement to supply the required services at a reasonable cost or demonstrate through independent audits and certifications (e.g., SOC 2) that the provider is sufficiently secure. Naturally, it also will be necessary to regularly perform an audit of the data center for compliance or obtain evidence of their current certification. Typically, this is not wanted by many data center providers since it carries an additional operating expense, but doing so could be a competitive advantage. The opportunity to obtain this information is greatest at the time of contracting since there is pressure to close a sale.

With physical facilities, one typically manages risk through contracts, certification, or mitigating controls. The occasional visit to the data center to perform a vulnerability assessment sometimes is necessary for due diligence. Before doing so, it is advisable to have a checklist of items to verify and a scheduled appointment; otherwise, the visit is going to consume a lot of time and be less productive. While the checks shown in Figure 10.3 are common for traditional data centers, just because infrastructure is provided as a service does not eliminate the need to validate them.

10.2.3 Infrastructure—Bare Metal

Inside a data center, one or more other service providers may offer bare metal services. These typically are rack mounted servers of varying sizes along with network switching and available routing. This approach to IT services removes the requirement for significant capital expenditure and tax chores such as depreciation schedules. With bare metal, the risk analyst is challenged with not only the physical data center security controls but also those of the systems and processes that enable installation and maintenance. Often forgotten are the risks associated with firmware updates, which could have an impact on operating systems and

Area to Check	Potential Vulnerabilities
Presence of video cameras at ingress and egress points (This should also include any loading docks.)	The ability to retrieve and examine video will help identify exploited weaknesses in the process, training, or technology of the physical security systems.
Secure doors requiring authentication for access	Some doors may have week, disabled or easily exploited controls. For example, a door propped open will allow unauthorized people to enter the facility. This can be mitigated with audible alarms when the door is open for too long.
Personnel on-site to confirm deliveries and visitors	Lack of oversight can lead to unauthorized visitor access to secure areas.
Escorted access to authorized areas	Unescorted visitors and delivery personnel can steal or damage infrastructure without accountability or control.
Register of visitors with arrival and departure times	This basic logging is fundamental in root cause analysis for security incidents.
Clear identification procedures for visitors—photo ID on official document	Failure to confirm identity of a visitor is a weakness that can allow abuse of identity and privileges at the responsibility of the impersonated individual.
Physical security controls on tenant cages	Large data centers have frequent visitors to the same floor to nearby cages. The potential for accidental or intentional access to the wrong cage can result in significant damage.
Power and network ingress to facility from diverse demarcation points	Destruction of power and network resources through accident can result. Redundant sources make it less likely to impact critical systems. Also, cameras in these locations can help detect or discover perpetrators.
Background check procedures for personnel	This basic control reduces the risk of hiring individuals with ulterior motives or who are more likely to succumb to temptation of former criminal associates.

Figure 10.3 Physical security vulnerability controls.

hypervisors. Government agencies have been known to tamper with the firmware installed in numerous systems. Furthermore, it is common that personnel are on staff who will provide "remote hands" assistance in installing and maintaining system components.

When employing Infrastructure as a Service (IaaS), patching of the operating system remains quite important and is often the purchaser's responsibility. Some IaaS providers supply low-level compute resources such as bare metal and a network with remote console access to the customer. The orchestration software then is used to automatically deploy the operating system of choice.

In other cases, the customer provides a copy of the operating system or purchases it through the provider and has it installed with the aforementioned remote hands service or uses custom software supported by an orchestration tool such as OpenStack. Once installed, scanning and patching of the OS are the customer's responsibility.

Consistent with our methodology, there are security questions concerning the implementation and maintenance of infrastructure services provided. A limited list of these should include:

1. Certification of the physical facilities—mentioned in the last section.
2. Background checks of personnel providing remote hands services.
3. Certification of the service provider and the scope of the services for which they are certified. Since there are network components, how are these governed and protected?
4. Licensing of any software provided to the customer should be confirmed. If this is not practical, the contractual control becomes more important in order to assign liability.
5. A process of verifying that the installed software version and patches provided are current or at the desired functional level for the planned application.
6. Service level agreements for patching of firmware should be an amendment to any contract with some penalties for failure.

As described earlier, there are many dimensions to infrastructure and one cannot be certain of how these dimensions are built. For example, with a bare metal server, it is well understood that there are physical connections to a "top of rack" or "end of row" switch. At a minimum, these connections carry the physical and data link layers in the OSI model. In scenarios where the service provider has many bare metal servers and electronically sets up one for your use, it is possible that the server is provided through a blade in a chassis of many bare metal servers. The top-of-rack switch may very well be integrated into the blade chassis.

At this point, the cloud consumer will have to rely on the practices and standards of the service provider. The assumption then is that connections to other servers of other customers in the same data center, rack, or chassis are kept segmented.

It is not always necessary to simply trust the service provider at this level. There are simple technical checks that can be performed in the link state of server connections. The MAC address list on the server and logs associated with locally established connections can be used to detect improper configuration. Furthermore, if privacy is a concern, and since it may be possible in a shared physical infrastructure to copy traffic from an interface, encryption can be used as a mitigating control. There are additional options in a bare metal setup related to virtualization. In fact, it is quite likely that the bare metal leased from a cloud service provider will have additional layers of virtualization in order to facilitate segmentation scalability across a global infrastructure. For example, it is entirely possible for two bare metal servers on the same virtual local area network (VLAN) to reside in two different physical data centers. This will become more apparent in the next section.

10.2.4 Infrastructure—Virtualization

Other flavors of infrastructure services include virtualization and migration, which enable deployment of the OS on various types of hypervisors using orchestration tools such as OpenStack supplemented by custom software. Similar problems arise except that now, more software components are introduced by the provider. These components generally are invisible to the consumer of virtualization services except as they are presented in custom software from the provider. Some of the provisioning services often provided are:

- Self-service on-demand computing (bare metal and hypervisors)
- Bare metal (configure BIOS and operating system installation)
- Network (create VLANs, virtual private networks [VPNs], virtual switching)
- Storage (build volumes, install databases, cryptographic services)

As will become apparent later, layers of virtualization often are used to aggregate the layer-2 or data-link layer to reduce the cabling and overlapping VLAN IDs that become exhausted at a certain volume of connections. For example, two different tenants in a single physical data center may have the same MAC address that, if part of the same VLAN, would cause a conflict. Furthermore, there may be the same VLAN number used by each of the different tenants. To allow massive scalability and distribution, the provider is able to form each of these VLANs into a virtualized VLAN transported over layer-3 (network layer). As a result, traffic out of the single chassis for multiple tenants is transported and aggregated over the same cable yet remains invisible to others. The mitigating controls for risk remain the same. At higher layers in this stack of connections, encryption can ensure confidentiality and integrity of the connections.

Once virtualization is provided, all of the underlying infrastructure is in place and the consumer of resources takes responsibility for the virtual machine where the operating system is installed. This necessarily leads to configuration of virtual

network elements including virtual routing and switching as well as optional virtual firewalls. The customer has considerable control without having to be concerned with the hypervisor and the underlying bare metal.

However, it is natural for service providers to offer services with high capacity components to add other services such as network firewalls, VPNs and storage. These add-ons spare the customer the responsibility for mundane administrative tasks and costly licenses that instead can be spread across many tenants.

With this level of resources provided, the security analyst should consider how each component will be validated for its security and compliance attributes. A firewall or VPN can be quickly validated if the administrative interface is available for inspection or configuration by the tenant. In other cases, look for the ordering process to define the characteristics of the service. For example, storage can be selected and should indicate clearly that encryption is used and specify the strength of the encryption as well as the permitted key sizes, types, and how keys are conveyed and/or stored for the customer.

10.2.5 Infrastructure (Limited Platform)—Containers

Open-source software is pervasive and quickly adopted, especially in the container realm. Many of the services obtained today are provided through containers and the microservices architecture they support. "Containerization" supports software delivery cycles that are very short, which is consistent with the Continuous Integration/Continuous Delivery (CICD) model in DevOps. There is a huge benefit in a fast-moving business world where new features and enhancements can be deployed very quickly.

Scanning and patching *application code* in a production container environment is not very practical. Containers are deployed as Docker™ images from a repository. Deployment is commonly orchestrated using Kubernetes, an open-source system for container management. Since each image that is deployed is a set of code components to perform a function, it is important to understand fully what is inside that container and any vulnerabilities that should be resolved before deploying as an executable.

Once in production, it is helpful to keep track of the last security state of the container. To do this, consider taking advantage of labels. These labels, which are structured as name:value pairs, can be used to readily identify the attributes of the particular pod release. Assuming that the reader has some familiarity with the Kubernetes architecture, an example of labels might be: "version=1.1." There are many strategies to using labels. For example, if a historical library of the releases is retained, it would be easy to determine that "version:1.1" is not current and 1.3 is the latest. With an accompanying record of vulnerability patches between those releases and the dependencies on other pods, a DevOps team can determine the risk in that particular module. Another effective approach is to note the ID of the last scan (e.g., "scanID:0123456"). By comparing this with the latest scan ID in your scanning software, it will be easy to identify the deltas.

There are commercial tools available to manage code releases and perform vulnerability assessments of images in container environments. It is even possible to perform more container-centric scans with such tools. Patch management may involve wholesale replacement of the system or virtual machine upon which the container environment is built. This has become a common practice since the deployed image is considered a stable, tested release and is not always patched in place. The principles of operation in the world of containers notwithstanding, process and policy are instrumental in successful patch and vulnerability management.

This section has discussed container services as a limited platform. The reason for this is that the service provides a choice in technologies that contribute to the overall ecosystem. Tools to manage and monitor containers are sometimes included along with add-on features such as firewalls and software development or deployment tools.

10.2.6 Platform as a Service (PaaS)

Platform as a Service (PaaS) is a category of combined services for a complete development, deployment, and operational solution. It includes all of the underlying infrastructure, although the details are not always clear to the customer. When a customer wishes to be very specific about the underlying IaaS components, the price likely will rise. But what the customer is often seeking is simplicity in a solution that does not involve maintenance of licenses, resolution of interoperability issues, and underlying capital costs. PaaS is well suited for the IT department that cannot afford the considerable dedicated expertise or is not prepared to make a long-term commitment to a specific technology.

Development platforms, on the other hand, provide additional services that will require close attention with the understanding that there is underlying risk that cannot as easily be identified in the software development lifecycle. The stack of services will raise more questions concerning how security is provided in each of the services:

- Are the code modules and development tools patched and at what frequency?
- Are there service level agreements (SLAs) associated with patching?
- Do the build and test tools include some means to test for code weaknesses?
- Does the solution include tools to track development progress and code submission?
- Does a test environment accurately reflect the state of the production environment?
- Are there sufficient tools to help meet industry-specific security requirements (e.g., encryption of data at rest and transaction logging)?

- Are data centers available in the required jurisdictions?
- Can secure, alternate channel connections be established for development apart from the production and testing environments?

Managing vulnerabilities in this model is, in some instances, a shared responsibility with the service provider. Application development, network controls, and directory infrastructure may be provided, but the purchaser of the service is responsible for configuration and monitoring. The buyer's responsibility not only involves discovery of vulnerabilities in the platform and development components but enforcing service level agreements for remediation of findings.

It is this last point that is important with PaaS. The service provider has complete responsibility for the operating system, physical host, and physical network. The contracts and service levels related to securing these components should be carefully considered.

One very popular hosting platform, Docker,™ provides an environment for storing, deploying, and operating cloud applications or components typically designed as microservices. For those unfamiliar with microservices, these small programs are intended to be small, quickly replaceable, and function as a part of one or more applications. One might think of them as the tradition subroutine that can operate independently. This application design enables rapid upgrades and replacements of these parts without affecting the entire application. Naturally, the operating flexibility facilitates the "DevOps" continuous improvement/continuous delivery (CICD) model.

Herein lies a specific challenge with vulnerability management principles. A container could have server components, including some commercial or open-source libraries. Each of the components of the container requires checks for vulnerabilities, remediation, and potential regression testing of the application prior to production deployment. The applications to be deployed in a container typically are kept in the previously discussed "trusted repository," which has restricted access and modification controls.

Once deemed ready for production, the code from these repositories can be deployed to the available and appropriate compute resources such as a pod on a Kubernetes cluster on a host. But, where and when are the vulnerability scans performed? Note: Even the Kubernetes people will tell you that it is not a full PaaS solution. However, it can be considered an integral part of a platform; therefore, I liberally cite it as an example.

Some service providers provide the ability to scan for vulnerabilities and deliver a report when the code is in the repository. It is also possible to perform the scan in the container that gets deployed, although that may be a bit late for remediation. This might be considered regression testing in a nonproduction instance. But both approaches serve different purposes. Scanning of code in the repository provides

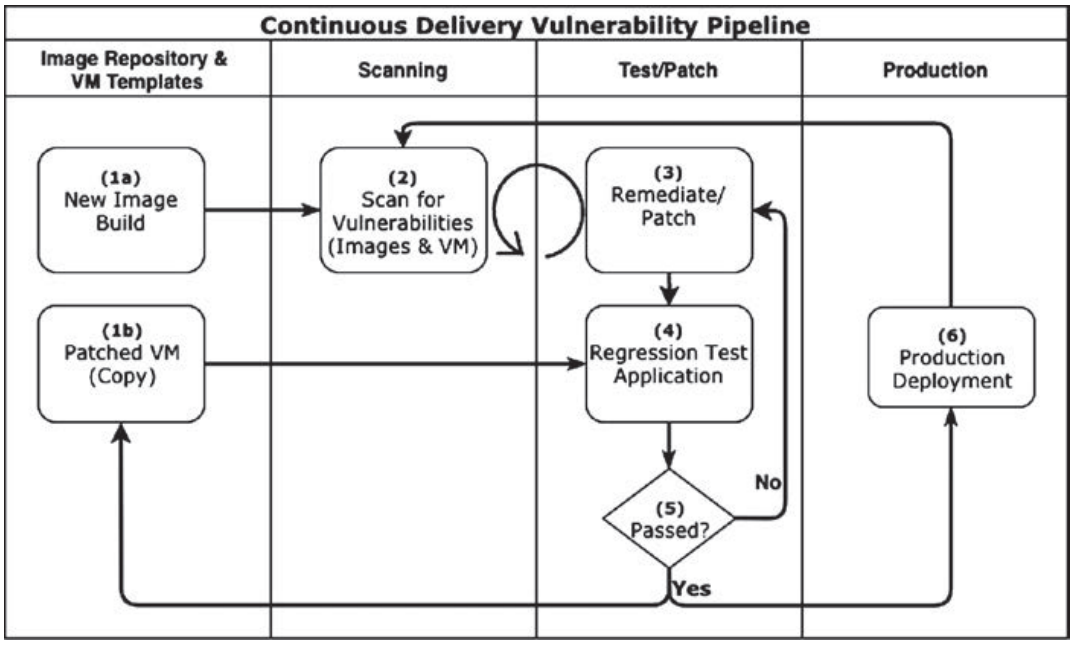

Figure 10.4 Continuous vulnerability management pipeline.

the current state of code that ready for deployment. Remediation activities can be performed in place and prior to production.

Scanning production containers has the benefit of providing a current running state. Some production containers have a low attack surface and can endure without vulnerabilities requiring immediate remediation. On the other hand, repositories can age or the code may change with new libraries required. The latter libraries may contain vulnerabilities requiring checks prior to production release.

These possibilities make it important to formulate a consistent process for scanning and remediation aligned with the development-deployment cycle. Figure 10.4 shows an example of a continuous scan, patch, and delivery cycle that accommodates a container environment hosted on a virtual machine. Both the VM and the containers are scanned for vulnerabilities and regression tested. Note that on the left side of the process, there is a container image repository and a "garage" of virtual machine templates. These contain the latest, patched and production-ready instances of a given solution. They can be deployed to the appropriate pod and cluster when needed.

With each release of the application, this process must be followed left to right:

1. Update or build:
   a. the code in the container
   b. an instance of the virtual machine
2. Scan vulnerabilities in the VM
3. Repeatedly submit containers into a repository until automatic scans are clean.

4. Once steps 2 and 3 are successful, perform a regression test of the application. Many development teams seem in such a hurry that they skip this important step.
5. After step 4 is complete, a copy of the VM should be retained in the aforementioned "stable" prior to being sent to production, step 6.

10.2.6.1 Other PaaS Risk Concerns

The development and delivery of applications is not the only function of a PaaS solution. There are important services typically available from a menu that can be added to any solution: Storage, Networks, Databases, Network and Host Security, domain name system (DNS) services, Load Balancing, and any number of other components that an organization might find useful. The simple idea is to get the customer interested in one thing and use that as a means of selling them additional add-ons that will entrench them with the provider.

This is not a bad state of affairs and, in some cases, is essential for a growing or even sprawling enterprise. However, there are strategic vulnerabilities involved with these choices. One might consider any systemic components that might present an unforeseen attack surface. For example, if a single vendor is responsible for all of the components to be used, there may be specific code libraries or architectural attributes with weaknesses yet unknown across the entire solution stack.

One way to mitigate this risk is to choose a provider, architecture, or solution set that has vendor diversity built in.

There are service providers who provide a low price and offer most of the services as a homegrown solution. This will create a greater requirement that the provider have strong development and risk management practices. Indeed, the homegrown solution may offer features and functionality most suitable to the consumer of services. Furthermore, the price can be compelling and the risk low.

Other providers have a well-curated collection of integrated tools to help manage risk for customers of their platform. For example, once you lease computing resources, there are add-on controls for network security, vulnerability scanning, container management, secure storage, x.509 certificate management, log collection and monitoring, and firewall and intrusion prevention. These services can quickly help manage vulnerabilities and compliance requirements. Specifically, IPS services function as virtual patch mechanisms to mitigate vulnerabilities in your solution for a period of time during which you can scan for the vulnerabilities and remediate them.

As for compliance with the reporting requirements, there are services to help directly. The European General Data Protection Regulations (GDPR) are partially met through encrypted database services, and some PCI requirements are met through firewalls, certificates services, and security monitoring solutions.

All of these services should be monitored and reevaluated for quality and relevance to your needs over time. It is not uncommon for a consumer of services to operate on autopilot and spend on services that are no longer needed. Additionally,

both security needs and service capabilities change. These need to be periodically realigned to avoid stealth gaps that can raise risk.

10.2.7 Software as a Service (SaaS)

From the point of view of the service consumer (customer), SaaS is a very simple approach for providing application services. None of the elements of the platform are exposed to the user. Instead, only the user interface of the applications is presented, which provides administrative simplicity. This also means that the hardware, software, network, storage, power, and data center are invisible. The provider does not wish to expose any of these details because lower cost, speed, and simplicity are the objectives.

To achieve these objectives, the service provider has to aggregate some of the resources—notably software licenses, hardware, storage, and network components such as routers, switches, and firewalls. This allows for economies of scale and better profit margins and/or reduced prices. None of these facts establishes any judgment of the security or compliance of the solution offered. A partial list of the areas of security not revealed to the customer directly include:

- Software development methodology
- Release schedule (sprint frequency)
- Model, version, and patch level of software, including:
  - Database engine
  - Operating system
  - Programming language
- Host monitoring
- Network resilience and load balancing methods/performance
- Firewalls, IPS, and Web application firewalls
- Virtualization tools
- Vulnerability scanners (infrastructure and application)

A typical agreement made with a provider is that they will provide these components and manage them appropriately according to best practices and that the customer need not be concerned. While this may be adequate for an inventory management system at a mid-size manufacturer, it certainly will not suffice for highly regulated industries such as banking and commercial finance. Taking the full range of customers and the potential purchase sizes of their services into consideration, there are several questions to consider:

- Is the provider financially accountable for their actions or inactions related to security?
- Are there service level agreements for performance and availability?
- How frequently are application vulnerabilities patched (read release cycle)?
- Are the developers trained in secure coding practices?
- Can application logs be provided in a continuous stream?

- Is there any security standard followed for which they will be audited?
- Can the application integrate with an identity management solution?
- Will you receive a copy of the audit report?
- What is the frequency of the audit?
- Will any findings of the audit be remediated according to a published plan?
- If there is a security breach, how soon will the provider notify you?
- How much information will be provided about the root cause and scope of the breach?
- Are there monetary consequences for the provider's failure to patch, detect, report, contain, or remediate?

For SaaS customers who spend considerable amounts of money for a custom SaaS offering, it is sometimes possible to negotiate more visibility and control the operating environment. While it is still unlikely that the provider will allow the customer to perform their own scanning and patching of the solution, it is practical to ask that certain scan and remediation service level agreements be established. It is also possible to expect that the state of some security controls be present. These are typical for a solution that is customized for the buyer's operational and regulatory needs. A simple example is the integration of authentication into the application through the corporate directory.

So, why not ask that critical vulnerabilities in the software stack be remediated within a certain time frame and that issues and outages that may impact the customer be reported in a timely fashion? Furthermore, as the buyer, you are entitled to know when there are events that may affect the service and what measures will be taken to remediate these. This is especially important when it comes to patch management. Even the most basic of web applications that are broadly used by the public will inform users when the service will be offline for maintenance or improvements.

Additionally, the DevOps operating model today can allow for frequent and rapid changes to the environment with very little loss of service. It is not uncommon for features, enhancements, and fixes to be deployed every 2 or 3 weeks. That in itself is a double-edged sword. Some of these rapid releases can introduce new vulnerabilities that end up getting fixed in the next CICD cycle.

10.2.8 Blurring the Lines—SDWAN

Software-defined networking is another cloud technology that blurs the line between a conventional network and a virtual one. This innovation allows the customer to implement their own WAN and integrate with their LAN without purchasing dedicated circuits or expensive services like MPLS (Multiprotocol Label Switching). Instead, the customer defines the behavior and features of their WAN through software. The network transport is conventional Internet access with connections among sites protected through a variety of options in VPN services. Decisions about performance, content caching, routing, and prioritization of traffic are configurable by the customer.

Another line that is blurred is between infrastructure and software (IaaS and SaaS). Although there commonly is a piece of hardware installed at the perimeter of the network at a given location, there is also software at play. This software is produced and maintained by an SDWAN (Software Defined Wide Area Network) manufacturer and sometimes installed or patched by a service provider (e.g., telecommunications company). So, vulnerability management decisions will have to be made just as they would when an SaaS provider is engaged.

The configurations and performance-tuning options are managed using one or more servers with specialized software. This software can be hosted and managed by a telecommunications carrier with expertise and technical resources, or it can be directly managed by an organization of sufficient size and resources. The logical extension of this technology is to further integrate the routing (control plane) and application (data plane) into a virtual or cloud software infrastructure. So now, the infrastructure that we customarily use for our applications and data can also support network routing and performance management.

However, similar to Software as a Service, this solution uses software and related configuration items to manage traffic flow. In many cases, the same questions concerning software solutions apply to SDWAN. The software must be built securely and the implementation and management should be properly governed. Assuming that an external service provider is responsible for management of your SDWAN implementation, Figure 10.5 shows some critical security questions to ask.

Technical Monitoring	Inspection	Contract	Risk/Vulnerability Topic
N	Y	Y	Are the personnel properly trained and their backgrounds checked?
N	Y	Y	Are the hosting facilities securely managed for all the physical controls discussed in the data center and infrastructure topics?
Y	Y	Y	What is the patch status and service level for the software?
N	Y	Y	Does patching follow a prioritized approach and have service levels?
N	N	Y	Will the provider inform customers when patches are being applied and if any operational impacts should be expected?
Y	N	Y	Are WAN connections encrypted and with what levels of encryption?
N	N	Y	Do the administrators have access to traffic flows through the network?
N	Y	Y	What are the controls on monitoring administrative activity consistent with other laws?
Y	Y	Y	Are there additional infrastructure or application services added on to the service?
N	Y	Y	What are the controls around the added services?
N	Y	Y	Are annual audits conducted on the operating controls?

Figure 10.5 SDWAN security questions.

10.3 Scanning and Remediation in Cloud

Up to this point, the discussion has been concerned with establishing the mindset of cloud and the vulnerabilities or broader risks therein. Now, if we consider cloud as a means of replacing infrastructure and platform but not full control of the application or environment, the need for vulnerability scans remains essential. We have touched upon, in part, the idea of checking for vulnerabilities in a dev-ops environment. This requires further elucidation.

There are two scenarios for vulnerability scanning that are most relevant in the present technology environment: virtual machines (VMs) running on a hosted physical infrastructure and containers running in VMs or a a service-provider container cluster (e.g., Docker™).

10.3.1 Scanning Virtual Machines

When a VM is hosted in a public cloud, you have little to no control of either the hardware or hypervisor. The network is invisible to you, and there is no clarity in the patch state of the underlying software. However, if you consider your hosted VM as your domain, it must be protected, updated, monitored, and governed. The vulnerability scan is a critical component of this and can be implemented in one of two ways: network scan and host agent scan.

10.3.1.1 Network Scanning a Virtual Machine

When scanning a VM over a network, it is important to take into consideration the real position of the scanner versus the target VM. Network scanning can take a heavy toll over the network components given that they are likely a combination of physical and virtual networks. The scanning of 65,000 transmission control protocol (TCP) ports (rarely needed) in a discovery phase will consume considerable resources and may not be permitted by a service provider.

However, it is possible that there is dedicated or at least sufficient physical infrastructure and a minimally shared hypervisor that may be able to accommodate such a scan without causing a denial of service. In all cases, however, it is advisable to minimize the impact of the scan by reducing the number of ports and IP addresses in a given period and to decrease the frequency of scans to a tolerable level. This is most feasible in server environments because there is more control over the services that are installed and the changes made to configurations.

One might also say that if two VMs are on the same subnet, they will not have a problem with the heavier network load. While this may seem logical in a physical environment, it is not at all so in the virtual world. In a cloud service provider, it is very common for a VM deployed on one host in a/24 subnet to be on a different host and possibly a different data center from another VM in the

same subnet. There may indeed be 5 to 10 physical and virtual network devices touched by a scan of the same subnet.

An alternative approach, which has gained rapidly in popularity, is to install an agent on the hosts. This solution is more elegant and has adapted and grown to be very stable since its introduction many years ago. Agents need not perform port scans and have a very low impact on compute resources. Furthermore, remote command and control of system resources can be gained through the agent and vulnerability information securely retrieved.

Local scanning agents have also become increasingly integrated with patch management. So, through a single agent, the security administrator can now monitor the patch status of the VM and initiate patches to system components manually or automatically. Detailed templates, rules, and policy enforcement of configuration items, patches, and event software installations are available.

Having discussed agents versus network scanning, I would like to call attention to a potential discrepancy with agents, the PCI standard, and your auditor. As of this writing, the PCI Data Security Standards call for a network scan of all 65,000+ ports of a host. That is a very big hurdle for a network-based scan, and yet an agent need not perform this action. The agent can determine immediately if the host has any of these ports open. In fact, some agents are able to monitor these ports from inside the host and prevent unauthorized ports from being opened. Furthermore, in a complex network environment, there are some devices such as small-network DSL modems that respond to a TCP port that has nothing to do with the server or application behind it. So, an external network scan may reveal open ports where in fact there are none on the PCI-scoped system.

If you choose to use an agent and have PCI-scoped systems, consult your auditor to reach an understanding of security versus blind compliance with a standard. The goal of PCI standards is to secure cardholder data, and there are many solutions. The goal of vulnerability scanning through agents is to secure the endpoint of any kind in furtherance of securing the overall cloud ecosystem.

10.3.1.2 Scanning Containers

As mentioned in the last section, containerization provides speed and short delivery cycles for new features. But, speed is a double-edged sword when it comes to risk. Although this is a benefit when remediating vulnerabilities, the patch windows are also shortened. One must take advantage of this situation and prioritize remediation of vulnerabilities in code equally with new feature development. There is a tendency to asymmetrically deploy features and related vulnerability patches. This can occur when a new feature uses a service that has a vulnerability not yet reported. Then, shortly after deploying the new feature, a critical vulnerability is discovered in one or more of the system components that are dependencies for the feature

(e.g., operating system libraries or add-on libraries for databases, encryption, or monitoring infrastructure).

While this is sometimes avoidable, the reporting of vulnerabilities, patch availability, and software deployments can create a situation where a minimum of 3 weeks, and possibly 6 to 9 weeks, may pass before sufficient regression testing can be performed. So, in the Scum methodology, the product owner will have to take responsibility for prioritizing the vulnerability into the backlog. Prior to this, it is very important that the vulnerability analyst or risk manager communicate the severity and impact potential of the vulnerability as early as possible.

Another concern is the vulnerabilities discovered through network scanning. The container space commonly is implemented in VMs that are rendered almost immutable so that the operations team knows the operating state of the environment. Do not be deceived by the term "immutable." Turning off the remote shell or removing the command shell altogether does not truly render a host immutable. There are other ways to attack a system.

Furthermore, traditional scans are not practical for a production container environment because the operating environment often has numerous containers with private, internal IP addresses. The container platform typically provides firewall rules to control access among the containers using numerous criteria. Privileges are assigned allowing certain containers to communicate with other containers using only specific protocols. Checking the vulnerability state of a VM over a physical network layer and trying to reach subnets internal to a specific host is not going to work well. Orchestration of a network scan from an outside scanner can become very complex. In all likelihood, the internal IP addresses of a container cluster have no external routes advertised to make them reachable.

Vulnerabilities must be detected in the build phase (pre-production) where a replacement virtual machine or system configuration is prepared. The goal is to have a continuous integration and delivery process that includes detection and remediation of code vulnerabilities as well as testing prior to submission into a trusted container registry. Figure 10.4 illustrates an example of such a process.

Fortunately, the market place has solutions for scanning a production container environment, and additional investment may be required. Any organization operating containers on a large scale should consider the advantage of a scanning agent installed on the host. Also, there are some scanners built into the container management environment that can analyze the binary of the installed components and report any significant vulnerabilities. A valuable selling factor to the development team is that the known state of the production environment can be better described with a scanning agent that can report all vulnerabilities on demand.

Equally useful is the ability to automatically perform scans anytime a new image is pushed into a repository. That approach works very well with a CICD (Continuous Improvement, Continuous Deployment) operation saving time and reducing the time gap between a change and vulnerability discovery. There will be more discussion on the process related to scanning containers later in the Platform as a Service section.

10.4 Cloud Risk Management Strategies

Since cloud services are by definition delivered by an external organization, the control of vulnerabilities and the ability to identify them is often limited. This happens for many reasons that the service provider does not always consider when originating their solutions. At first, the provider sees the natural monetary benefits to their customers but does not fully share the same risk-management vision. This is not an attempt to deceive but rather the natural evolution of the market.

As a result, it is important that the consumer of these services take measures to mitigate the risks inherent in a rapidly evolving technology ecosystem. The consumer is king when two service providers must bid for services, and the one meeting the most security requirements will gain an advantage. Some very useful instruments for managing service provider risk are: contractual agreements, implementation of mitigating controls, a vendor-risk assessment program, and a competitive bidding process in a diverse market.

10.4.1 Service Diversity

When feasible, consider the option of building diversity in the choice of service providers. The simple example is where there are multiple applications to host that require no direct communications. These can be spread appropriately across providers or data centers within a single provider. The fact that multiple suitable service providers are available can be used as a means of lowering the costs of services while tactically obtaining the mix of security controls desired.

Another scenario is to identify for each application what controls are essential and which ones are optional and at what addition cost. Using this information, one can select the optimal service providers having an acceptable balance of performance, capabilities, risk and cost. Since wide area networks can be virtualized at higher layers of the OSI (Open System Interconnect) model, one need not be bound to a single provider.

10.4.2 Contractual Agreement

Contracts are the legal tools that are the risk managers' friends because they outline a general framework or scope for security governance where the service provider is concerned. Some security managers consider contracts a complex and arduous

process involving attorneys, conference calls, and the exchange of many document revisions. It behooves the risk manager to align the risk appetite with contracts and learn the art of mitigating controls. Armed with this information, work with your organization's attorneys to ensure that your risk concerns are addressed in the contract. Prior to contracting, however, establish and follow a vendor risk assessment program to identify the appropriate controls.

10.4.3 Vendor Risk Assessment Program

Vendor risk assessment is a common practice in any organization. For IT migration or ongoing operation in cloud, there is no special exception. The methods have to evolve but the general practices are the same:

- Define the business purpose and solution
- Build threat models where applicable
- Profile the technical environment
- Determine the risk
- Build governance requirements
- Establish controls and metrics

Some IT organizations are mistrustful and believe that only they can truly manage risk and give comfort to the organization. That approach is incredibly expensive and is not necessarily effective when trying to flexibly meet rapid changing business requirements. One approach to managing risk in the cloud is to determine the real risk to the organization and gain an understanding of the data, functions, and real exposure. Just because the application is moving a cloud model does not mean risk increases. It is only the insecurity experienced from a false perception of loss of control. To address this, consider the following guidance:

1. Become informed—Gather detailed information about the solution that meets the business requirements. Security details *are* business requirements. Virtualization of an entire application (lift and shift methodology) can provide easy assurance but is not always the best solution. Perhaps it is a stepping stone to containerization, which can increase resiliency and performance.
2. Establish external assessment program—Audits, attestations, penetration tests, and scans all have a place in vulnerability management for cloud. Know which combinations of these are truly necessary to identify well-defined risks. Simply demanding a control without purpose is often a waste of time and resources.
3. Build a cloud monitoring platform using controls that match the present and future. If the organization is accustomed to managing CPUs, memory, patch status, security events, and incidents, it should also be able to continue this with added controls. For example, rather than collecting simple syslog events from cloud (many cloud solutions no longer produce syslog), use tools such as

Prometheus. This is a an extremely versatile and efficient means of monitoring using time-series name/value pairs. It provides consistency and versatility to data collection more efficiently than trying to parse syslog events.

10.4.4 Compensating Controls

If the contractual process is not the right route and you do not control the risk environment, you will have to find compensating controls. Consider these worst-case circumstances:

- A vulnerability goes unpatched in a vendor system.
- The entire application (SaaS) has an operational vulnerability that is likely to remain present for several months.
- One or more of these vulnerabilities are determined to be severe and unacceptable to the organization.

Under some combination of these circumstances, it will be necessary to find a compensating control. Scanning, patching, contracts, and moving your data to another provider are not practical. You can take other actions such as:

- Implement a technical control to:
  - Hide the critical data from the application through encryption.
  - Restrict access to the cloud service by forcing users to pass through your security/content inspection infrastructure.
  - Use an external risk monitoring service that can support your particular web application.
- Implement one or more administrative controls (in many cases, you should do this anyway):
  - Require users to reaffirm their access to the application more frequently.
  - Perform monthly or quarterly data quality, access, and usage reviews.
  - Restrict the user roles permitted in the application (reduce threat level).

10.5 Cloud Security Assessments

The industry working groups have established some tools to help with automation and consistency in security assessments. Center for Internet Security (CIS) Benchmarks are used to assess the configuration of cloud IaaS services using best practices established in coordination with the CIS and numerous security professionals (Figure 10.6). Some cloud service providers support automation through vulnerability scanning solutions to use application program interfaces (APIs) to access configuration data. This information can be used to assess a service's compliance with the benchmark. It does not, provide a complete picture of the vulnerability

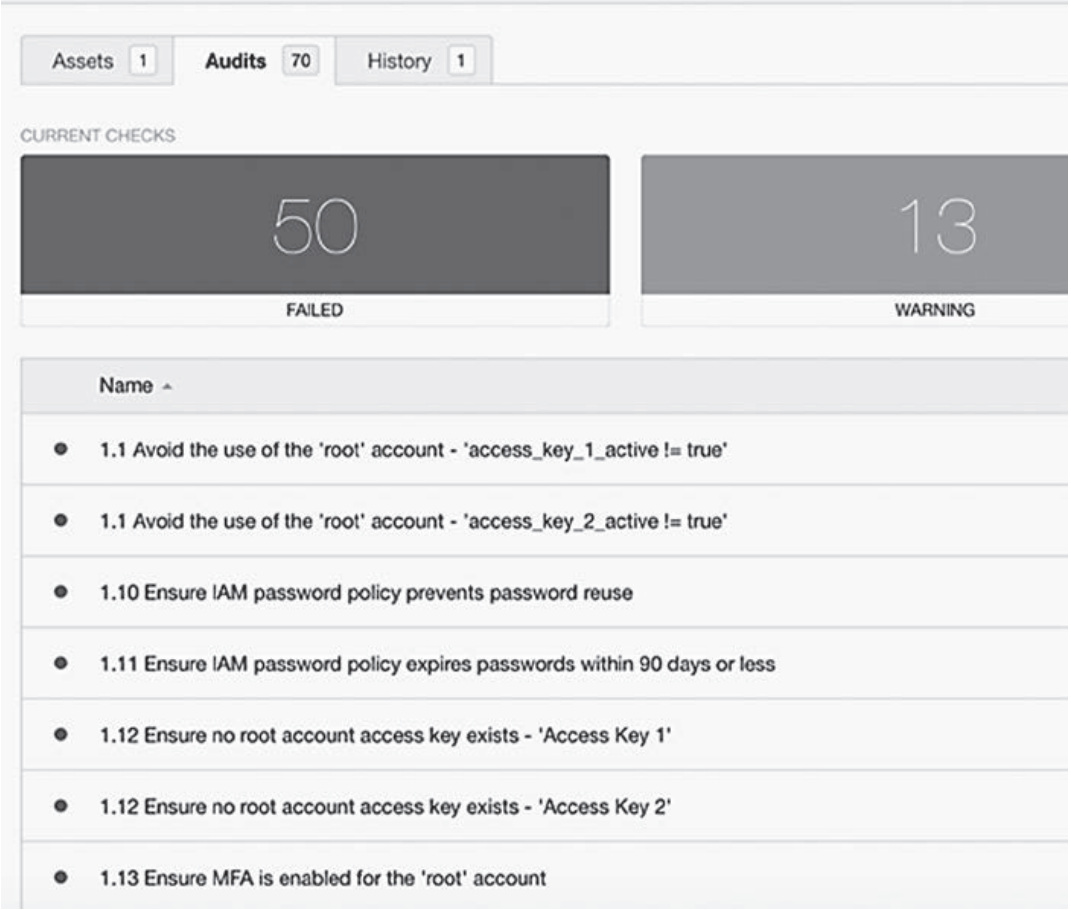

Figure 10.6 Cloud provider security compliance checks.

state of a cloud service. One simply gets assurance of the best-practices configuration consensus found in the benchmark.

Another useful tool could be the CloudTrust Protocol (CTP)[4]. This protocol calls for transparency to support "digital trust." The simple idea is to provide a certain level of transparency into the security state of the cloud service offering. The provider is not necessarily going to provide you with vulnerability scan results nor will they provide any configuration items. Instead, the CTP consumer of a cloud service gets high-level information about the compliance and security state of the service. A commonly cited example is up-time of the service, which is conveyed in an SLA-compliance formation composed of promised state versus observed state. For example, up-time over 10 days was at least 99.8% and the provider commitment is 99.7%; thus the consumer knows there was compliance. In order for this protocol to work properly, the service provider is obligated to establish service and asset classes. These are used to map specific attributes of a service or asset in order to accurately convey compliance.

In general, CTP is a major improvement on transparency to cloud service consumers. However, with so many layers in a cloud service (as described earlier in this chapter), it is unlikely that a large service provider will be able to provide accurate and consistent information for all components. On the other hand, CTP provides consistent methods of obtaining compliance and risk information should the service provider find ways to make it reliable. Naturally, in a multi-cloud provider environment, the metrics from one provider may not be the same as those from another provider. For example, provider A may supply availability information for a server but not for a network. Another provider may provide network availability but not for a server.

10.6 Conclusion

In all cases of cloud services, it is the responsibility of the consumer to gather the required information before, during, and after using cloud services. This translates to contracting, operating, and terminating services. Security and assurance tools and technologies are available in varying and inconsistent forms. Small consumers will be able to insist on minimal capabilities by "voting with the feet." They will select the provider that has the best reputation, price, and most acceptable offering. Large consumers of cloud will be able to negotiate a bit more and, in many cases, build their own security into the solution. That is the great thing about cloud. A vulnerability at one layer can be mitigated by strength at another layer where the consumer has more control.

End Notes

1. https://www.nist.gov/publications
2. Competitive Enterprise Institute, https://cei.org/i-pencil
3. European Union Agency for Network and Information Security (ENISA), 2009
4. Cloud Security Alliance, 2018 https://cloudsecurityalliance.org/working-groups/cloudtrust-protocol/#_overview

CHAPTER

4

CRISIS MANAGEMENT AND DISASTER RECOVERY

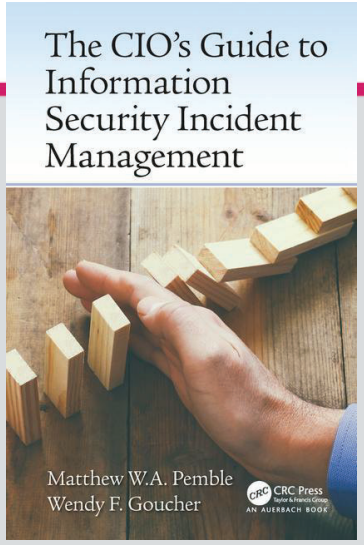

This chapter is excerpted from

The CIO's Guide to Information Security Incident Management

by Matthew William Arthur Pemble, Wendy Fiona Goucher

Learn more

22

Crisis Management and Disaster Recovery

And this mess is so big
And so deep and so tall,
We cannot pick it up.
There is no way at all!

Dr. Seuss
The Cat in the Hat

Information in this chapter:

- Types of crisis
- The role of the IR manager in a crisis
- The team
- Disaster recovery

Introduction

As has been said in Chapter 5, the difference between an incident and a crisis is a matter only of degree and when applied in a business environment, who needs to be involved. However, some regard it as a more subjective matter; "One man's crisis is another man's challenge." This may resonate well with positive thinking gurus, but it is an unhelpful way of looking at any ongoing business critical situation because it can mean that that senior-level involvement is not incorporated early enough. This can result in the situation becoming more of a crisis for longer than it needs to be.

Sometimes security IR may find itself part of a larger IR situation. The two key times this will happen are in terms of a disaster recovery situation where there has been a failure at scale of part of the business

or its supporting services, or in crisis management where an incident of any cause is having such an impact to the business that it is being run by as senior executive team.

In either case, there will be differences depending on the extent to which IS is part of the cause of the problem or you are merely supporting an incident with a very different cause. For example, a disaster recovery situation resulting from a ransom-ware outbreak will involve you in a very different situation than would be caused by denial of access to premises as a result of a chemical pollution problem.

However, there can be many similarities: the requirement for confident contact with senior members of the organization who may have little or no understanding of the art of the possible and the potentially extensive contact with external bodies such as the media or regulators which will require additional skills and place different working priorities on the IR team as a whole.

There Are a Number of Types or Stages of Crisis

Fundamental Crisis: From the first moment, it is clear that this is a situation of sufficient seriousness that senior management need to be kept informed. In pure business continuity terms, the crisis that resulted from the physical attack on the Twin Towers, leaving aside the human effects or the effects due to the political motivations for the attack, was a fundamental crisis from the moment the first plane hit the first tower. Further, knock on effects from the incident were revealed, but the situation was already a crisis.

Revealed Crisis: This is where the full extent of the problem is not apparent until the incident is being investigated. The nature of the incident does not change, but the full extent of the situation is not clear on first analysis. Following the email trails in an investigation into a mainframe outage, of itself a very significant incident, uncovered a range of serious misconduct ("gross misconduct" in UK employment law terms) amongst hundreds of staff members, including several issues that required to be reported to law enforcement.

Developing Crisis: Subtly different from the revealed crisis, this is where the impact of the issues under investigation increase during the investigation. A phishing attack may be detected after initially targeting a small group of users, but the attackers may release

subsequent attack blocks affecting a much larger population. Equally, an initial malware attack might contain a programming error rendering it merely irritating rather than effective and this may be corrected in a second release.

Imposed Crisis: This is an incident that would not, in other circumstances, be a crisis. However, because of other factors such as the public profile of those involved, whether people or organizations affected, the situation needs to be handled with reference to the organization's executive. A phishing attack which overtly targeted high-profile public figures may need to be treated differently at the customer and media management level, even if the underlying technical investigation and corrective actions are identical to those involving the general public.

The Role of Incident Management in a Crisis

Whatever the cause of a crisis, the most valuable asset in successful handling is information. Without information even being able to describe the shape or seriousness of the crisis can be difficult. Until these parameters begin to be defined, gaining control is impossible. Two of the key roles on incident response are to help to gather information about the incident and work to re-establish control of events.

Crisis Management

Where the impact of any incident to the business is such that it poses a significant risk to ongoing business activities or even an existential threat to the organization, good corporate governance requires the executive to play an active part in the management of the problem. As an example, the very large fines introduced with the General Data Protection Regulation, up to 4% of worldwide turnover, are likely to cause material issues for even the most well-funded business. Which, one suspects, was the intention. Equally, allegation or suspicions of gross misconduct against members of the executive themselves or key personnel, while less financially awe-inspiring, may still threaten market confidence or critical business areas or programs. There should not be much surprise if sudden changes in the business focus or direction that you were given emerge. As well as the IR reports, the

executive will be receiving guidance from other specialist areas such as legal, regulatory and media relations and that advice may make certain aspects suddenly urgent or even significantly more important than the IR team might have been working under.

Expectancy Management

It has already been mentioned this with regard to reporting upwards and sideways in less serious incidents, but it in a crisis situation it is vital. It is reasonably likely that the executive do not have a glowing opinion of the business competence of IT to which the IR function may be associated, and their understanding of the limits of IR technological capabilities (particularly if funding for that shiny new monitoring or investigations equipment was not forthcoming) will be poor and their tolerance for technical jargon negligible.

Especially where there is a mistrust of IT, let alone IR, it is essential that they have a common and consistent reporting channel which should be the IR manager, in which case that role alone is bound to be a distraction from the active management of the incident. A preferred solution, which is one Matthew has used and recommended in a range of different situations, is that the executive have the IR manager's immediate boss as their contact person. If that is not possible, depending on the size of the organization, it may be possible to bring in a second incident manager or, failing that, rely on an experienced lead technician to take most of the active management burden from you.

Consistency and Simplicity

In order to facilitate this, it is helpful to have a standard reporting template, which should be the same as or very similar to the one you use for slightly less serious incidents.

It should be clear and simple, and lay out the current state of the issue, any actions or decisions you require from them, and provide, in so far as possible, the necessary information for those decisions to be taken. It is essential that reports are provided at the time they are expected, even if this means that they are less complete than ideal. Even if there is a major change in the incident behaviour, it is

essential to get a holding report off. The timings that the executives have committed to with regard to their obligations to inform other key stake holders may not have been made clear and an absence of information when the executive need it may hamper the quality of the information they have to pass on.

Media

Nobody should attempt to deal with the media unless they have been properly trained to do so. Ideally, the IR function should not deal with them directly at all; it is best left to corporate communications or media relations specialists, whether internal or external. It is a very good idea to have a set of template forms pre-prepared, discussed and practiced with media relations. These can then easily be populated with the necessary information for the particular event. The most critical thing is never to put any media spokesperson in the position of saying something that is either untrue or that could be, even if true, robustly challenged as misleading or evasive. In particular, underestimating the scale or impact of an issue is unwise. It is worth remembering that particularly with the expansion of locally focused media, modern publishing has meant that an issue can appear much more critical through the lens of different media organizations, especially with their own priorities and prejudices, than it seems to you from the corporate perspective.

The Team

While the team are concentrating on their analysis and reporting, it is important to remember that the communications channel down to those working on the incident, especially where this involved the communication of any changes in priorities or required activities, must be effectively maintained. In both exercises and live incidents, we have seen situations where to an observer it appeared that were two totally different incidents underway, one being competently managed by the blue team and the other well under the control of the white team. Indeed, on the worst occasions, neither bearing much relation to what a third team, the red team, were actually doing. Depending on your background, the requirement to be the hub of a many faceted

communications network may be something that you are entirely happy with or something that you need to develop through training and exercise.

Part of the role of the IR manager is to protect the team and its morale from any sudden changes in direction. Obviously, it is not always possible to insulate them from these changes, but if requirements are properly explained and the value of the work the team had been doing up until that point is not undermined, this should minimise the impact on morale. This can mean the manager being in the position of taking a degree of flack from both sides, but that is just one of the perks of the job.

Obviously, if there are significant additional demands from investigation or analysis activities it may be necessary, if practical, to call in additional resource so that the core work, whatever the incident is, can be maintained. Depending on the resources available to you, this may not be immediately available, so careful planning and a degree of honest and considered anticipation will be helpful.

Disaster Recovery

An organization should have a disaster recovery management framework or plan, probably integrated with the lower level business continuity management framework. Those in charge of the IR function need to study this, talk to the responsible practitioners and ensure that the higher levels of your security incident management mesh seamlessly with the disaster recovery plan. Remember that the team can be involved because of security causes or impact of security actions or simply because it is perceived as responsible adults who can valuably assist at a time when is "all hands to the pump." It is also essential the teams' contact details are kept up to date in the business continuity call-out lists.

Revealing the Case Studies

Amber Inc.

Amber Inc., as the larger company of the two here, has the greater opportunity for executives who would normally not be involved in standard IR scenarios, needing to be kept aware of crisis situations.

Also, there may be incidents which fundamentally relate to the physical, non-technical parts of the business, but which impacts areas that are covered by the remit of the IR team. For example, where a catastrophic weather event affects the supply of basic utilities, such as electricity, this may seem to be an access issue for staff, but it can also affect the provision of IT services. Where the organization is larger, it may be that those working in disaster recovery and incident response don't regularly meet. This may not because they don't want to; it may be because they are based on different sites and, in this case, possibly even in different countries. However, there will almost certainly be DR situations which need to have IR input so Amber Inc.'s DL and IR responsible managers should attempt to agree on the approach to situations where IR may be able to contribute positively to the solution.

While co-ordination may sound like such a sensible approach, it would not be contested, there is, in larger organizations like Amber Inc., an additional problem. There is a greater probability of competing egos between IR and DR. Why should one of the team change process in order to suit the process of the other? Especially where both perceive their team as successful, taking a step away from the successful approach could be to reduce the positive career impact for the senior manager. There may also be misunderstandings between the teams or it may be difficult to agree on the prioritisation of elements of an incident. In the worst-case scenario, both teams may blame the other for mistakes and attempt to "score points" for their lack of fault.

The potential for significant reduction in the effectiveness of management of a situation requiring input from both teams means that it is vital that steps are taken to attempt to mediate, and ultimately achieve, effective co-ordination.

Jade Ltd.

For Jade Ltd., the damage to reputation and core operations that could result from an ill-managed crisis should help to motivate staff to endeavour to work together effectively. This being a smaller organization there are fewer places to hide mistakes or work-arounds that the DR or IR team may use. There may be a good reason for their evasion; it may be, for example, that the budget that was allotted to them was insufficient for the process they had hoped to use. However,

with such a small organization, especially when the DL and IR teams have to work together, these work-arounds will need to be mutually understood, as well as working for both teams.

The need to have a pre-organised and practiced approach to media queries is essential in this case. Should a significant incident happen on an otherwise slow news day, then any fault in the interaction with the media, which includes control of official social media output, could mean that the event receives greater coverage than expected. This would mean that any absence of clear information could be filled with conjecture. News abhors a vacuum and journalists are well-practiced in the art of filling them – and generally not in favour of those who didn't answer their calls. Ensuring that such queries are answered effectively should be a core part of the creation of disaster and incident response processes for any organization commonly dealing with sensitive information, whether or not that data is involved in the incident.

CHAPTER

5

ENSURING A CONTINUOUSLY CYBER-RESILIENT ORGANIZATION

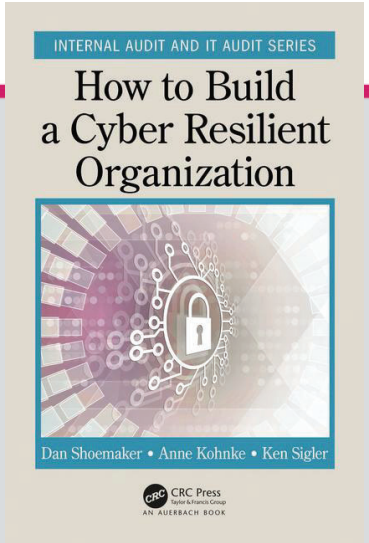

This chapter is excerpted from

How to Build a Cyber-Resilient Organization

by Dan Shoemaker, Anne Kohnke, Ken Sigler

Learn more

Chapter 8

Ensuring a Continuously Cyber-Resilient Organization

At the end of this chapter, the reader will understand

1. how to create and sustain a long-term cyber resilience program;
2. the standard elements of sustainment architecture;
3. tailoring out functional elements of noncritical asset recovery from standards;
4. the assumptions associated with sustainment architecture;
5. the purposes and importance of documentation and archiving;
6. the goals and success factors of the audit process;
7. the steps involved in ensuring long-term sustainment.

It All Starts with Infrastructure

This chapter presents and discusses the detailed principles that underlie the development of dependable, large-scale organizational infrastructures. Those principles are rooted in strategic process definition and planning. The purpose of both of those comprehensive processes is to establish continuing and consistent day-to-day functional cyber resilience operations.

Specifically, this chapter will detail a practical organizational process for developing and maintaining a stable, complete, and sustainable cyber resilience response. The goal of cyber resilience is to ease the impact of disruptive events on

the information asset base. This assurance is guaranteed by utilizing well-defined organizational control processes to protect critical and noncritical information assets and services. In real-world terms, the necessary trust is underwritten by the creation and sustainment of an operational control infrastructure for cyber resilience.

The goal of this formal structure is to foster and maintain a complete and consistent environment for the conduct of the everyday cyber resilience process. The infrastructure makes the abstract elements of cyber resilience real and tangible to the people in the organization. In order to do this, the infrastructure formally documents and embodies all the relevant organization-level strategic and operational concerns that might be required to successfully ensure critical asset protection. In addition to protection, the infrastructure will also embody and sustain an effective and reliable noncritical asset recovery mechanism.

Infrastructure development also embodies the practical steps that are required to ensure the continuous improvement of the conventional cyber resilience process. An organization's cyber resilience infrastructure ensures suitable execution of the everyday practices needed to ensure protection of critical assets and recovery of noncritical ones.

Organizations maintain strategic governance mechanisms as a means of ensuring that the fundamental processes and practices that comprise the cyber resilience infrastructure have been fully and correctly aligned to the particular organizational purpose. Governance also ensures that those processes are executed in such a way that they ensure the long-term, effective achievement of the overall business goal of full and complete protection of their information assets. The infrastructure makes that abstract notion real and substantive, in the sense that it oversees and regulates the day-to-day operational activities that perform the cyber resilience process.

Logically, the specification of those activities must be both detailed for any given organizational application as well as strategically viable as part of the overall conduct of the business case. These divergent goals are satisfied by a well-defined and formal array of deliberately interrelated operational processes that are capable of identifying, analyzing, responding to, escalating, and learning from all adverse events.

Embedding the Cyber Resilience Process in Day-to-Day Business

In order to be effective, the processes that comprise a cyber-resilient infrastructure have to be performed as part of the overall day-to-day operation of the business. Thus, cyber resilience must be embedded in the conventional business as an everyday operational process. That is the case because every form of behavior that is designed to mitigate risk must be continuously functioning in order to ensure

against the exploitation of vulnerabilities. The real-world behaviors that comprise that continuous protection are generically termed persistent or organizational "controls."

Persistent control is an important concept in architecture because it represents the practical behaviors or actions to be performed. The control universe is by necessity diverse and complex. That is because the threats and attacks they are tasked to address are often intricate and multifaceted. More importantly, they are not always computer based. Exploits like social engineering and insider theft center on the human factor rather than the machine. That implies that it must be possible for companies to oversee and manage an increasingly complex set of day-to-day practices that must be reliably executed by every member of the organization. Given the layers of complexity associated with something as abstract as virtual work, this is a particularly difficult task.

Therefore, the creation of a reliable cyber resilience architecture requires a wide variety of precisely targeted behavioral processes and automated controls. The real-world realization of these controls must be

1. planned based on intimate knowledge of the threat environment;
2. fully integrated to produce a synergistic result;
3. economically feasible.

Most of today's global business is carried out in distributed, multilayered, multivendor, and even multicultural environments. All that work must be satisfactorily coordinated and controlled in order for the company's stakeholders to trust the outcomes. That coordination and control are often very difficult to achieve given today's complicated business challenges.

Therefore, oversight and control of complex work are built around the assurance of a common and stable infrastructure of management controls. A control infrastructure helps the organization guarantee uniform and consistent outcomes from all of the many operational components of a multifaceted corporate environment. That common point of control is embodied in the everyday control architecture of the organization.

Given the inherent complexity in control formulation, the only way to ensure a systematic architectural solution is through a single rational development strategy. The strategy must ensure the trustworthy, long-term cyber resilience capability of the organization. And, it must address all likely threats to all items of critical and noncritical assets. A solution such as that must be systematically designed to merge all requisite behavioral and technical controls into a single, fully coordinated operational process across the organization.

The architecture must be implemented as a substantive element of organizational functioning and then evolved as a way to meet the changing needs of the environment. The generic term for this type of large-scale infrastructure management process is "governance." Because it becomes harder to ensure reliable outcomes

as business organizations grow and diversify, a rationally planned governance architecture of realistically targeted controls is important for any form of successful vulnerability management.

Throughout this text, we have approached the topic of cyber resilience as if it is a strategic governance function. The purpose of adopting that stance is to avoid a piecemeal solution. Whether they are formally documented or not, every enterprise is managed by a set of commonly accepted practices. In most corporations, those practices represent an individual unit's or manager's understanding of the proper way to carry out a specific task, and they tend to embed themselves in the organization over time.

Typically, every organization's cybersecurity practice has evolved this way. That is, organizations have developed solutions bit by bit as the situation arose rather than as an across-the-board strategic governance process. The problem is that a piecemeal protection architecture doesn't work in a world of persistent threats. That is because the operational security response will not fully embody or address all the concerns for the asset base.

The alternative is to formally define, design, and install a comprehensive array of management controls, which are aimed specifically at optimizing the effectiveness of the security function. This is done as a coordinated entity. As with any complex deployment, these controls can only be substantiated through a rational and explicit planning process. The practical management term that is used to describe a process of designing and deploying a specifically targeted universal infrastructure response is "security architecture." That term effectively describes the strategic approach and purposes of the cyber resilience development process. Moreover, it is this strategic architectural concept that drives the deployment of the types of controls we have been discussing in this text.

Security Architecture

We need to spend some time explaining the general principles embodied in security architecture. The term "security architecture" was coined to describe the strategic function that underwrites proper due diligence in the assurance of the organization's information assets. Security architecture consciously builds an intentional structure of rational elements and interorganizational relationships that will sufficiently ensure the business' information assets.

A security architecture approach centers on the creation of a comprehensive organizational control system and security culture rather than building separate individualized security solutions for each problem. Essentially, the business defines a coherent framework that embodies all the necessary roles, their associated behaviors and practices, and the strategic policies necessary to achieve a single coherent, organization-wide cyber-resilient solution. There are five processes that underwrite the development of a security architectural framework.

Scope

The first of these principles is scope. In essence, scope addresses protection issues in the light of resource constraints. In the worst case, that might involve leaving a nonessential or low risk resource outside the boundaries of the protection scheme.

Proper scoping entails the execution of a deliberate process to establish the boundaries of the solution. In some respects, it is the most critical step of all, since the extent of the territory that must be covered by the security solution will dictate the form and extent of the rest of the process, and in the real world that extent implies a conscious balancing between need and the resources to meet that need.

Accordingly, the underlying issue that scoping addresses is: "How does the organization get optimum assurance out of its resource investment?" In day-to-day practice, this means that it must be possible to make an informed and intelligent decision about the level of risk that can be accepted within the resources available. Obviously, anything can be secured if enough money is thrown at it. But no organization has the money to effectively put a cop on every street corner. So, a deliberate process has to be undertaken that balances deployment of the assurance response against the likelihood and material consequences of the threat. Factors that might enter into this process include considerations such as: What is the level of criticality for each information asset and what is the degree of assurance required for each?

During cyber resilience prioritization (Chapter 4) that determination is captured on a ten-point asset classification rating scale that ranges from "not needed" on one end of the scale all the way up to "the business would close without this" on the other end. This can be used to support the tough decisions that will have to be made at the intersection between the decision to protect or leave out the defense.

Because technology evolves, the process of establishing a cyber resilience scheme is always dynamic. For decision makers, proper scoping maximizes resource utilization and provides the foundation for the rest of the process. In case of scope, this means that the protection modules are subject to ongoing refinement. That refinement is based on information feedback from other activities, particularly the risk assessment.

In essence then, although the first step is to define the boundaries of the modules that will protect the critical assets, these are not fixed. They are always subject to change as the realities of the environmental circumstance dictate. However, if that decision is made based on precise knowledge about an established perimeter, this knowledge is obtained by formal assessment.

Standard Risk Assessment

Assessment is the second essential element of the cyber resilience infrastructure development process. That is because assessment is a primary player in any form of risk-based security system development. The assessment function ensures that the organization fully understands the risks inherent in its security environment and the associated set of requirements.

Since risks don't come with convenient labels, criteria are required to assess risks. Those criteria are normally embodied in the form of a requisite best-practice reference model, which documents all the necessary controls of a standard governance model. In that sense, the detailed recommendations of a governance model comprise the basis for risk assessment.

That best-practice model should be capable of optimizing the security governance process for that organization and represent expert consensus. Thus, most commonly accepted best-practice standards are documented as a logical structure within a domain and process framework. The model itself should allow an organization to evaluate the business risks and assign control behaviors to a common set of best-practice functions. By definition, a control objective is a precise statement of the desired result or purpose to be achieved by implementing a given control procedure for a particular activity. The standard that is utilized to make that decision should provide comprehensive criteria to make decisions about the overall completeness and correctness of a managerial control system for a given set of organizational assets.

In that respect, there are several reference models that can be selected. However, the Information System Audit and Control Association's (ISACA) control objectives for information and related technologies (COBIT) standard, the International Standards Association's (ISO) ISO 27000, and the National Institute of Standards and Technology's (NIST) NIST 800-53 are arguably the most frequently referenced models.

Any one of these frameworks can be used to judge whether a security infrastructure and its constituent activities are complete and correct. Whatever the source, the reference framework should specify and populate a well-defined and detailed organizational best-practice architecture that allows the organization to develop its own explicit control policies, practices, and procedures. In their practical application, COBIT and ISO 27000 are primarily oriented toward conventional business. NIST 800-53 satisfies the requirements of the Federal Information Security Management Act (FISMA) and therefore it is almost exclusively used in government.

All these frameworks start from one simple and pragmatic assumption: The operation must be managed by means of a set of logically related everyday behaviors, which taken as a whole constitute a complete set of security best practices for an organization. ISO 27000 has 14 process areas, the COBIT Framework defines four general process areas and a set of 34 high-level control objectives, and NIST 800-53 embodies 17 process areas. These are assumed to describe and embody all aspects of information and information technology (IT) functioning. By satisfying the specified requirements of each area, the manager can ensure that a capable IT control system is in place.

The actual control system is tailored top-down to any desired level of detail for any given organizational application. A standard tailoring approach is particularly important in the creation of security infrastructures, since the implementation of

the actual security system is always different in its particulars. A standardized security control infrastructure, which is defined at a commonly accepted, correct level of detail provides the communal policy and procedure reference points that are essential for the reliable definition of a practical security solution for a given application. Or put more simply, because a standard framework is stable, it can be trusted as the basis for making the rational trade-offs and adjustments that are always implicitly necessary to formulate a specific cyber resilience infrastructure solution.

Why is a well-defined and commonly accepted standard-based approach effective? First, it provides a continuous point of reference to reflect and accommodate the full range of difference in the security needs of the pertinent organizations. The ability to deal effectively with a wide range of potential applications is an essential quality in any standard-based security design process. That is because the threat environment is so wide ranging and diverse. Persistent threats can be identified at any time and from any source. So, the security response has to be creative and multifaceted. Since every potential application of the standard is different, the implementation process has to accommodate that range of difference.

However, since it is clear that there is a need for flexibility in the tailoring of controls from a standard, there are also good reasons for standardization. Standardizing the controls on a model of best practice ensures the time-tested effectiveness of the solution. Plus, with standard control areas, each implementation of the specific standard solution can improve the process through lessons learned. Standardization provides a standardized basis for measurement and metrics. And finally, in the practical world of business, it is naive to suggest that unique security solution should be developed ground up for every organizational application. So, some sort of standardized approach to the problem is also an absolute prerequisite for effective use of security resources.

In order to be at their most effective, security infrastructure development processes should have both maximum flexibility and standard structure, which might seem like a contradiction. The simple resolution is to create the architecture of the solution top-down from the highest possible level of concept. The largest and most comprehensive view of the solution can then be used as a general classification structure within which effective controls for each of the specific security areas can be addressed. Using this approach, a specific cyber-resilient control architecture can be constructed for any given project at any given level of definition inside the model best-practice framework.

However, the classification structure itself must always be consistent. That is, the overall control recommendations of the general framework model must provide the uniform best-practice elements and structural relationships that will allow a specific cyber-resilient control architecture to be refined to any desired level of detail for every project. In theory, using the approach, an optimum process can be defined for a given project.

The outcome of the specific tailoring of controls to satisfy each of the general security categories of the standard should be an operational process model, that

embodies best practice, and which specifically ensures cyber resilience for a given organizational application. Practically speaking, a well-defined and documented practical infrastructure of controls that is developed from one of the control frameworks mentioned above is of considerable value to the business. That is because that documentation makes the overall cyber-resilient control architecture and its attendant procedures tangible, and communicates them to the employees of the organization.

The behavior of any form of control system has to be documented. Documentation is an absolute ongoing requirement. In many respects, the documentation of control outputs is the absolute bedrock on which oversight and management rest. So, it would not be too much of a stretch to conclude that it is the documentation process that ensures that the cyber resilience function is properly run.

In effect, if the organization ever hopes to oversee and refine its control processes, it has to be able to keep track of what those control processes and activities are doing. Consequently, the documentation process has to be planned as part of the overall architecture of the cyber resilience process. That part of the plan articulates a strategy for recording and maintaining all the relevant output of a given cyber resilience control process and/or activity.

Conceptually, the documented realization of a standard model of best practice is at the other end of the application spectrum from the eventual cyber-resilient architectural solution. Because of the requisite tailoring, the solution itself is always organization specific, and generally not transportable to any other organizational setting. However, because it is derived from the standard, the tailored solution DOES represent the profession's most commonly understood best-practice solution to security control assurance in a particular case.

Building the Practical Infrastructure

The cyber resilience concept is based on two essential economic assumptions. First, organizational information is valuable, therefore any loss of information represents a loss of value. Second, the protection of information must be designed and implemented in such a way that the greatest value is obtained for the money invested.

The fact is that you cannot secure everything. Consequently, decisions about what to protect and how to protect it serve as the fundamental criteria in the development of cyber-protection for any organization. Consequently, any given architectural design process has to be capable of arraying a well-defined and functionally similar collection of components into an optimally cost-efficient infrastructure while leaving out functions that are not adequately justified by the value they ensure.

Infrastructure is a generic term that describes a real and applied set of systematic actions that are been put in place to ensure an abstract characteristic such as "cyber resilience." In essence, infrastructures comprise the necessary technology,

people, and operational practices to achieve some planned day-to-day operational purpose. The key terms here are "applied" and "real." An infrastructure framework establishes and substantiates practical, rational, and systematic everyday strategic purpose based on a specific investment of real resources. The investment is designed to achieve a defined goal. Each element of the infrastructure has a discrete reason-to-be, and each contributes differently to the ultimate goal, which in this case is a cyber-resilient organization.

Practical infrastructure design involves identifying and documenting all the requisite elements of the eventual solution and then specifying how they interrelate. Any technical, human, or practice behavior that does not achieve the desired organizational purpose would be a waste of money. Therefore, the process of developing the right mix of elements requires the organization to strike a rational balance between the outcomes of the proposed infrastructure component and its associated cost. That process is architectural in the sense that actual structures of related elements are being assembled based on some inherent logic. Therefore, the process of assembling an organizational infrastructure can be compared to, and has the same practical purpose as the architectural considerations that go into designing a complex building.

The end result of the cyber resilience control infrastructure development process is a coherent and fully integrated cyber-resilient architecture. This architecture is never a one-size-fits-all proposition. Instead, the determination of each element of the system must be driven by assessment and the degree of assurance required within a particular setting. And, the rule that underlies that condition is that all situations are different. Therefore, a correctly developed control response is always precisely tailored to correspond to the exact requirements of any given circumstance.

The process requires that there be some kind of standard point of reference to coordinate all security activities that will apply throughout the life cycle, which is the reason why a standard model of expert best practice is required. Cyber resilience control infrastructures are normally based on a commonly accepted, expert, standard model of best practice because there is no science to dictate the correct set of operating procedures.

Instead organizations have to rely on the industry's most up-to-date understanding of the control actions that work. The recommendations embodied in a best-practice standard are meant to serve as the basis for characterizing a wide range of security control functions; in concept, the recommendations of a best-practice standard are not standalone elements. They actually substantiate numerous facets of the security function that taken together constitute aggregate best practice.

Accordingly, as a set, the recommendations of a best-practice standard comprise a complete and tightly integrated control infrastructure solution that when properly configured produces a fully secure situation. In addition, best-practice standards provide a formal mechanism that allows the organization to explicitly specify the rules for control behavior as well as for how ongoing control performance will be assessed. There are eight commonly recognized factors, shown in Figure 8.1, which

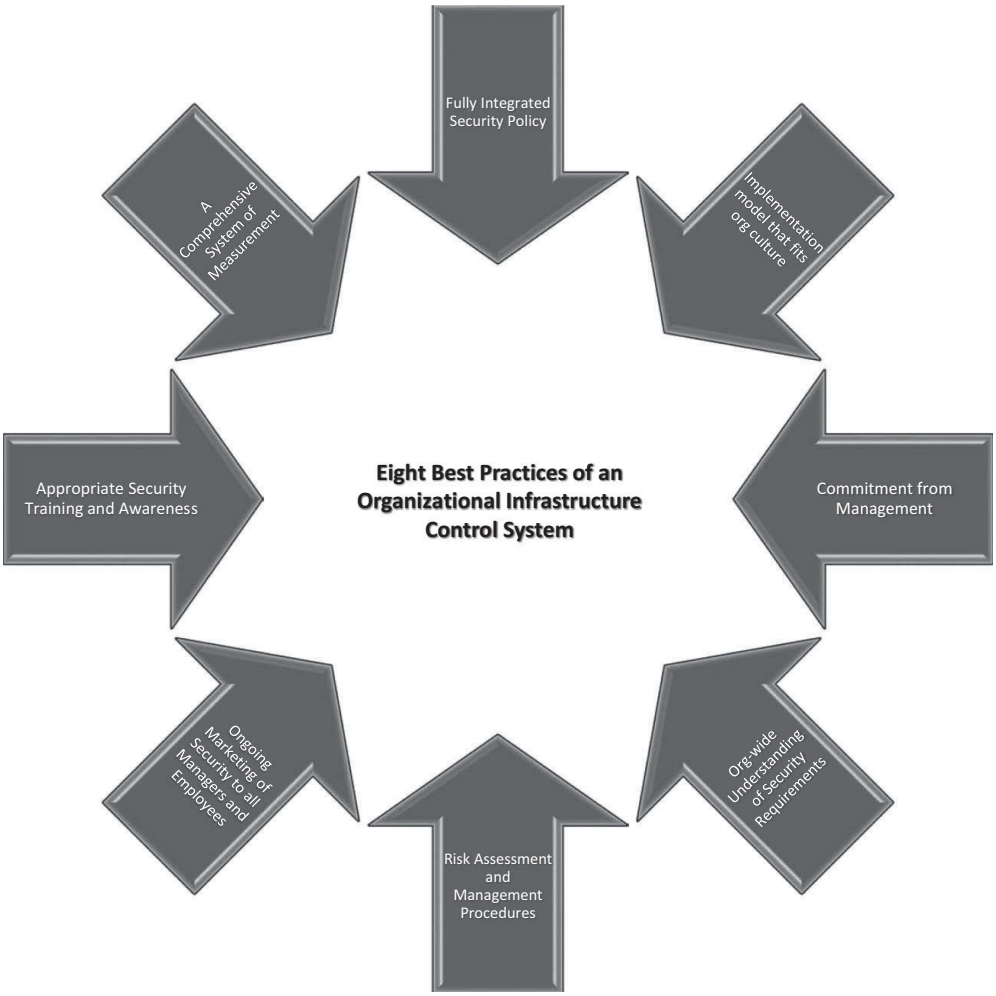

Figure 8.1 Eight best practices of an organizational infrastructure control system.

are considered to be instrumental to the overall success of an organizational infrastructure control system. These are:

1. *The existence of a fully integrated and coordinated, formally documented security policy*: Organizations are governed by their policies. Every policy is operationalized by a targeted set of controls. These controls must be complete, correct, and accomplish the purposes they are set to achieve. The proof that is taking place is confirmed by evidence. Audit is the function that is responsible for collecting, interpreting, and reporting that evidence.
2. *An implementation model that fits with organizational culture*: Human factors are important to the success of almost every element of security. But security behavior is hard to guarantee since it is governed by interpersonal rather than logical rules. Therefore, it is important to ensure that all rules of behavior are defined properly and accompanied by the appropriate motivational approaches. An implementation model that the employees hate will never be successful.

3. *Visible support and commitment from management*: The employees can be expected to resist the imposition of any form of security control. That is because security best practice is generally inconvenient. It makes people do things like enter passwords, when they would rather dive right in. The security people, and even IT administration, don't have the organizational reach or clout to enforce accountability. That can only come from the top. So, the cooperation and commitment of the people in the C-Suite is an integral part of a successful cyber resilience control process.
4. *Organization-wide understanding of security requirements*: Along with the commitment at the top, it is also necessary to make sure that everybody in the organization knows what the rules of behavior are and specifically what their accountabilities are when it comes to security. Understanding implies the need for the design and implementation of formal awareness, training and education programs that are aimed specifically at helping the employees of the organization understand exactly where they fit in and what is required of them.
5. *Formal risk assessment and risk management procedures*: Risk assessment doesn't end when the infrastructure controls are put in place. The risk environment is constantly changing and evolving, and the organization has to evolve its control infrastructure to keep pace. That is the reason why it is essential to have formal risk management embedded in the day-to-day operation as a means of staying on top of the risk picture.
6. *Ongoing marketing of security to all managers and employees*: Security is tiresome, so its necessity has to be explained to the people who do the organization's work. This is an ongoing operational responsibility that is key to sustaining an effective response. That is because people forget over time, and they tend to let inconvenient things slide. So, it is necessary to occasionally remind everybody why they have to do those irritating little security things.
7. *Appropriate training and awareness*: Managers and workers have to be appropriately and specifically trained in their security duties. That is because awareness and training programs make employees aware of the importance of technical and procedural controls. Thus, training and awareness are formal parts of infrastructure control system management. Employees should have scheduled periodic refresher training to assure that they continue to understand and abide by the applicable conditions. Therefore, some form of awareness and training for employees should always be undertaken as part of the overall cyber resilience process.
8. *A comprehensive system of measurement*: Infrastructure controls are monitored and assured by data that is generated from a comprehensive testing and review process. That process involves a carefully planned and targeted set of test, reviews, and audits that are designed to generate substantive data that will allow managers to understand the functioning of the control system at any point in time and for any given situation.

The Detailed Cyber Resilience Control System

A cyber-resilient architecture embeds a set of interrelated or interacting control elements in the organization in order to direct and control how its cyber resilience objectives are achieved. This is accomplished through a process-based cyber resilience control system, which is a network of many interrelated and interconnected activity elements. In that respect, a cyber resilience system is no different than any other complex system. Each activity uses the business' resources to convert some form of input stimulus into a predicted output. Logically, the output of one network action becomes the input of another action.

Therefore, the control activities that take place within a cyber resilience system embody a complex network of input-output relationships, which taken as an integrated whole comprise the single, process-based, cyber resilience assurance solution. To be functionally correct, the solution must specify an exact set of desired outcomes. These specified outcomes allow the organization to determine whether a specified protection goal has been achieved. This includes the specification of the level of validation that is required to authenticate that a given set of results meets the explicit criteria as well as any non-standard, post-task conditions that might be specific requirements of that particular activity.

The purpose of any type of practical architecture is to serve as a foundation for performing a large-scale, real-world business function. And so, in the simplest terms, the architecture of a cyber-resilient process provides the basis for ensuring the long-term security of the organization as a whole. It should be made clear here that architecture is not a management process per se, in the sense that it serves as the cyber resilience architecture. It is not the actual day-to-day, nuts-and-bolts execution of the process. Instead, the infrastructure represents a logical framework within which the work is done. However, to work properly, this structure has to be explicitly stated and its elements have to be explicitly related.

The standard architectural model that provides the basis for developing the applied system must encompass and describe the entire structure of the solution, from top to bottom. Therefore, it is necessary to define every process and practice at a level of detail sufficient to ensure proper operation of the cyber resilience solution. All the details of the elements that are used to categorize that solution must be expressly traceable and derivable from the higher-level elements of the architecture.

Development of a fully defined architecture would be difficult, if not impossible, if it weren't for the fact that standardized process models always contain common features that can be classified into a single category of basic operation. For instance, planning and documentation tasks represent a common set of requirements across most organizations. Yet, these all exhibit pretty much the same entry/task/exit (ETX) practices. Consequently, standard types of activities, like planning or documentation, can be understood and described as a common set of well-defined behaviors. Then, during the actual implementation process that common

set can be implemented and interconnected in various ways to achieve the unique purposes of the project.

For instance, planning has a distinct outcome, which is a substantive plan. So, planning always involves the same logical behaviors: information gathering, problem analysis, formulation of a substantive direction, documentation of that direction, development of an attendant monitoring process, and implementation of the plan. Organizations might do some of these individual tasks differently. But the entire set of tasks is still generally executed in that order, and they all are essential components of the process for developing a plan. Thus, there is no need to think through a new set of actions when a planning activity is required. All that is required is to customize the same standard actions to the new situation.

That is the reason why the basic unit of an architectural process model is called a task cell. In practical security implementations, the task cell is also known as a "control." Task cells are unitary functions. That is, each task cell is specified as the means to carry out one specific task, and one task only. The actions the cell defines have logical entry conditions that are required for proper task initiation. These include the inputs from any prior activities at all levels of abstraction. They also must produce an intended outcome from those inputs. The presence of the prescribed outcome serves as the basis for judging whether a task has been carried out properly. There are many task cells in any category of desired outcome in an architecture. The actual tasks themselves, that is, the behaviors that the cell carries out, are defined by describing a specific set of discrete actions that the cell must carry out in order to achieve a commonly understood and well-defined outcome for a given purpose.

These cells or controls are normally tailored to fit the stipulations of a standard model of best practice. That model stipulates the goals that support decisions about how these cells can be interconnected to achieve a higher-level purpose. The value and application of these models will be discussed later in this chapter. But there are other potential sources of association that can be tapped to dictate the possible interrelationships of a task cell. These include: (1) current standard operating procedures within the organization, (2) current or commonly recognized industry methods, and finally (3) any contract stipulations.

Constructing the Process Model for a Particular Application

Once a complete set of standard process cells has been defined, a process model can be constructed for a particular application. This is done by interconnecting the basic set of task cells in various ways to produce a tailored best-practice infrastructure of controls that will satisfy the needs of a given application. The idea is to incorporate into the solution only those behaviors that address the identified issues,

and produce the desired outcomes for that particular organizational purpose. In the real world, that involves approaching the solution in three different ways:

1. *A Staged View*: The problem is approached in defined stages.
2. *An Organizational View*: The problem is approached through a model of best practice.
3. *A Control View*: The problem is approached bottom-up through a highly integrated control set applied individually to each requirement, and aggregated into a system.

One of the most important assumptions of this text is that every organization can, and should, implement a cyber resilience process through a formal process, which has been appropriately tailored to fit the requirements of the organization's particular threat environment. The practical mechanism for ensuring this requires the following five steps:

Step One: A standard model of best practice must be selected to define the general form of the infrastructure's architecture, and its constituent controls adopted for tailoring out a complete solution. We have mentioned three popular models in this chapter. However, there are a range of commonly accepted, standard approaches to choose from. None of these is more or less valid in their overall applicability. It simply depends on the threat environment that the adopting organization faces. Obviously, a top-secret organization will adopt the standard model that the government requires. Whereas, it is more likely that a small business will adopt something less demanding, like the COBIT standard. The selection process is governed by the business case, more than it is something based on hard-and-fast rules.

Step Two: Once a standard model has been adopted, a set of behavioral task cells is tailored out of the general activity recommendations. Each behavior must formally operationalize a single control objective within the selected framework.

Step Three: The precise specification of the required input and anticipated output behaviors is done. This allows the organization to monitor and track the behavior of each cell.

Step Four: The specific ETX criteria and expectations must be specified for each cell. This is essential for monitoring ongoing performance.

Step Five: The requisite monitoring is described in an assessment plan. The standard and systematic assessment, measurement, and reporting that are described in that plan are then carried out on a systematic operational basis.

The overall alignment of the tasks in the cyber resilience architecture must be explicitly linked to the recommendations of the standard model of best practice. Yet, because every component in the architecture is unique, due to the tailoring for each situation, the practical solution also reflects the individual requirements of

each organizational application. In addition, the resultant cyber resilience architecture cannot be too rigid. The architectural solution has to be capable of modification as a way of reflecting the dynamic changes that are likely to occur in the threat environment over time.

Making Data-Based Decisions about Performance

Besides the need to assess where an organization is in relation to common best practice, using some form of standardized criterion, there is also the requirement that the organization ensure proper and effective decision-making about the operation of their cyber resilience system on an ongoing basis. This implies the need for a formal and systematic measurement and evaluation process that will allow the business to oversee the ongoing execution of the process using objective data. In general, the measurement and evaluation process must address the following practical concerns in order to obtain that data:

- *Operational Performance of the System*—What are the indicators of proper functioning?
- *Asset Prioritization*—Which assets are critical versus noncritical? How is this changing?
- *Risk Acceptance*—What are the risks of recovering versus protecting assets?
- *Benchmarking*—Are we achieving standard best practice?

In order to answer those questions, the cyber resilience measurement and evaluation process must have a system in place to provide objectively derived assessment data. Specifically, there must be a systematic means to ensure the monitoring and control of the evolution of the cyber resilience process over time. The ongoing understanding of the explicit operational behavior of the formal array of cyber resilience controls is derived from a continuous monitoring process. That process will allow the organization to benchmark control set performance over time. Ideally, management will be able to continuously evaluate the performance of its cyber resilience control infrastructure, and then take substantive action to ensure its effective operation in real time.

The measurement and evaluation process needs to be enforced by explicit, objective assessment measures that will document whether the controls that have been deployed in the cyber resilience system have met the requisite criteria for proper performance. That objective data should provide a basis for management to determine the adequacy of its current operational control set. To achieve that understanding, the measurement and evaluation processes should address the following four operational concerns:

1. Are the cyber resilience controls meeting the business' stated practice specifications?

2. Does the control set continue to achieve generic best-practice recommendations?
3. What are the residual risks of operating the system as it is currently configured?
4. Is the cost of executing the control set justified by value obtained?

The general requirement for systematic execution of the measurement and evaluation process is made operational by a set of well-defined and commonly understood control specifications that can be used to reliably and repeatably benchmark the everyday performance of the organization's real-world control practices. These assessments are done to obtain the data required to confirm control status on a systematic basis. The assessments utilize the outputs of what can be assumed to be proper control performance, which are documented in the standard for best practice that the organization has chosen. The business must perform regular operational assessments using these specifications.

First comes the need to precisely define and document the outcomes that must be obtained in order to indicate proper performance. There are four logical activities that are involved in the evaluation of objective outcomes. The first is the definition of explicit goals. The actions required to achieve those goals must produce objectively quantifiable outputs that can be documented and logged.

Second, the specific behaviors that will be used to characterize those actions and which will be documented for analysis must be described in sufficient detail to ensure that their presence or absence can be unambiguously confirmed.

Third, the specific behaviors that must be present in order to confirm that an activity is being carried out must be specified. That includes their timing and interrelationship. Finally, the prescribed corrective action required, should a behavior not be performed correctly, must be specified and documented. This data must be in objective terms, which can be operationally documented, audited, and verified as correct.

Once the desired outcomes have been itemized in detail, practical, operational measurement and evaluation processes must be designed and deployed to monitor and log those outcomes. Critical success factors that might be documented in this log include:

- Factors that document the satisfaction of specifically linked policies
- Factors that document organizational level functions, such as management
- Factors that document proper execution of the cyber resilience process
- Factors that document proper execution of an explicitly identified control
- Factors that document desired technical outcomes

The traditional assessment process is based on the presence of defined and documented outcomes and clear accountabilities. Intangibles such as s strong support/commitment of management, appropriate and effective lines of communication, and consistent measurement practices can also be factored into the assessment process as long as the outcomes can be observed and documented in objective terms.

All these factors belong to the business, and they should always provide unambiguous data that will allow the activity to be substantively measured. Ideally, each factor will be described in terms that will allow the organizational entity that is tasked with the measurement responsibility to be able to determine if, or when, an assessment process is successfully complete.

The practical assessment is performed in precise measurement driven terms. Depending on the outcome of that assessment, there might be a long period of trade-offs and refinement before an eventual decision can be reached about the effectiveness of the cyber resilience function. However, the final documented solution must objectively demonstrate that it addresses the strategic goals and business objectives of the organization.

Implementation Planning

Implementation planning constitutes the next practical step. The outcome is an appropriate set of controls. The security controls are embodied in a comprehensive, practical business architecture directly and verifiably addresses each issue identified in the risk assessment. The design of that architecture is always a creative, conceptual exercise in the sense that its final product is the "blueprint" of the cyber resilience control system that will eventually be put into place. These designs are rarely limited to technical documentation. In case of an overall cyber resilience architecture, the design is generally an overall architectural plan encapsulated in a policy and procedure manual.

All designs exhibit common characteristics. They are complete in the sense that they encompass an entire architectural solution. They are correct in that all elements of the solution, which logically should be present, are there. They are understandable in that they unambiguously communicate the form of the solution. Finally, all the elements of the infrastructure must be traceable to the standard model of best practice that was utilized to shape the architecture.

The design document that encapsulates all these characteristics will establish the quantitative foundations for the measurement function. But in addition, that document must explicitly embody the desired qualitative elements. These elements are subjective rather than objective. Nonetheless, they are useful in judging the success of the design itself. For instance, the stipulation that the system must be "reliable" is an example of a qualitative element.

In addition, the design must provide a clear direction for the integration principle, which is the next step in the process. In that respect, to promote the most efficient interaction, all the technical elements of the proposed solution must be integrated both at their internal interfaces as well as with the elements of the existing operation. Any subjective business process element has the same requirement.

Also, as the name implies, there must be some sort of well-defined and documented planning outcome. All the necessary relationship issues must be identified and resolved in the plan. And, from a resource management standpoint, all risk and

control issues have to be coordinated by a planned process. Because a plan provides the explicit direction for the infrastructure development process, it is the essential end product here. In case of a small business, this might be a relatively trivial documentation item, a memo of agreement for instance. However, where the security solution is either very large or complex, there is an implicit element of long-range planning, and the result is always a detailed plan.

In that case, the organization performs the activities associated with a normal strategic planning process. That includes rationalizing the solution against business goals and objectives as well as the formulation of a schedule and a contract. This justification process actively and intentionally aligns the overall form of the cyber resilience response to the organization's existing needs.

Control Integration

Control integration is the largest infrastructure development activity, in terms of the actual time and resources required to execute it. This activity might be just as appropriately termed "realization" or "implementation" because the eventual outcome of the process is the substantive cyber resilience system response. In case of cyber resilience, the general goal of the control system is to ensure authorized access to critical and noncritical assets and make certain that the processing and storage of information within the organization involves a well-defined and effective system of coordinated actions. This type of assurance is typically embodied through five common security attributes, shown in Figure 8.2.

1. *Authentication*: where an individual, an organization, or a computer proves its identity
2. *Authorization*: ensuring access to a specifically constrained collection of assets
3. *Confidentiality*: the ability to maintain the secrecy of assets that must remain private
4. *Integrity*: the assurance of the correctness and accuracy of an asset
5. *Non-repudiation of origin*: All communications are genuine and trustworthy.

The construction of the control infrastructure to assure this passes through eight logical stages. These stages force the business to think through the form of the cyber resilience architecture for that organization. It also ensures that the solution is both optimally resource efficient as well as continuously improving.

Assigning Investment Priorities

The first step in the process logically involves the gathering of all the information necessary to carry out the characterization and prioritization functions that are part of the standard cyber resilience development process. These have already been

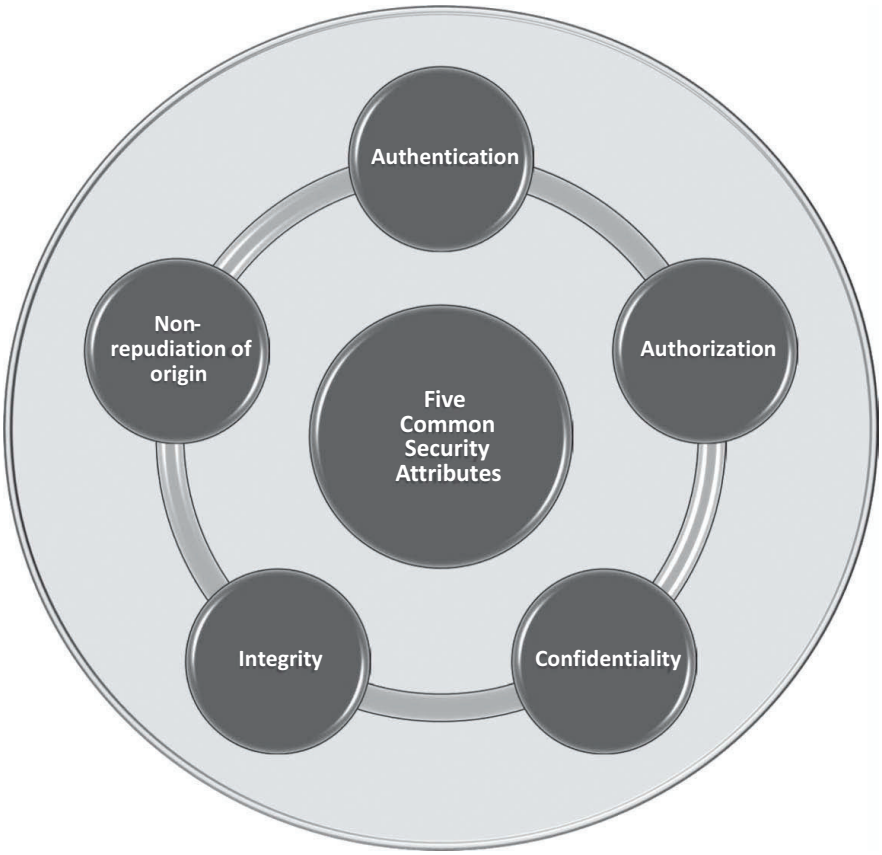

Figure 8.2 Five common security attributes.

extensively discussed in this text. Essentially, that process involves the identification, labeling, and valuation of each and every one of the organization's information assets, and the designation of the critical and noncritical asset baselines. In many respects, this part of the process resembles the asset identification activity that would drive any fiscal accounting or physical inventory process. The individual baselines of critical and noncritical items are differentiated and documented in a formal, organizationally standard ledger. Because that ledger is virtual, it is maintained under the dictates of rigorous, classic configuration management.

Once all of the organization's information assets have been identified and arranged into two coherent baselines, the next step is to determine their priorities. The purpose of this activity is to determine the exact security requirements of each of the individual items in the baseline. It is here that the important decisions are made with respect to the substantive controls that will be assigned to each asset item.

This obviously involves the risk assessment process, which we have discussed, and which is necessary to characterize all the direct threats, vulnerabilities, and weaknesses for each critical baseline item. The outcome of this data gathering process is that all the factors that might impact security for each item are understood and factored into the specific control response.

When the entire set of information assets and their related controls has been identified, the organization's decision makers do a rational assessment to establish the criticality of each item. This step weighs the value of the asset and the implications of all relevant threats against the actions required to address them. In pragmatic terms, this means that the potential impacts of the threat on the business are evaluated, and a level of importance or significance is assigned to the investment.

This activity applies to any of the identified asset item, and it requires consideration of every known threat as well as all other substantive strategic or business issues. This is where the trade-offs actually happen. The working principle here is that the practical impact of loss or harm to a given asset must be balanced against the resources that would have to be invested to ensure it's protected at the requisite level of assurance.

The eventual outcome of this process is a sliding-scale representation of all assets that is used to draw the line between the critical and the noncritical items. The scale ranges from deploying every control required to absolutely ensure the protection of a given asset, down to "the investment in controls isn't worth the value of the asset." This information is then aggregated into a total picture of the investment that must be made to ensure the resilience of all the critical and noncritical assets in the organization.

Since it is likely that the overall investment that is required will be greater than the resources available, this picture will have to be adjusted through negotiations. There will be some assets that will clearly belong in the critical category. Conversely, there will also be assets that will clearly not justify the investment. The art of cyber resilience planning is in navigating the grey area between critical and noncritical. This is where the real value of this part of the cyber resilience infrastructure development process lies.

Decision makers must learn to perform a triage that draws a clear investment policy line. They must prioritize the entire set of asset value and investment concerns in order to ensure that the maximum number of critical functions are guaranteed protected within the resources available. The functioning everyday cyber resilience control infrastructure is the eventual outcome of this process.

Rolling Out the Solution to the Stakeholders

Once the general cyber resilience plan has been developed and approved by the organization's stakeholders, the next step is to customize the real-world set of control practices and procedures. This is the point where the cyber resilience infrastructure is packaged and socialized to the entire organization as a substantive operational process.

The rollout formally establishes the cyber resilience control function itself. It embodies the substance of the cyber resilience process as an everyday reality. It

aids the organization's managers in resourcing for the long-haul control deployment and sustainment process. The successful implementation of a functioning standards-based infrastructure depends on the ability to communicate the requirement for a given set of actions to the target organization. Consequently, the rollout of a real-world operational control infrastructure from standards requires a series of focusing steps.

The purpose of these steps is to ensure that the behaviors that will be implemented are fully and completely understood by all stakeholders, and that those stakeholders are capable of fulfilling the purpose and intent of each of the generic recommendations of the standard. To do this, the implementers have to communicate both a high-level understanding of the process, as expressed in the standard, as well as a specific set of control behaviors that implement those recommendations.

The practical realization of the company's cyber resilience process is built around an exact specification of the details of the control infrastructure. In that specification, the requisite control behaviors must be explicitly understood and executed. In order to measure and assure this, the expected behaviors must be executed at a level of unambiguous performance. Therefore, the precise set of actions that must be taken to ensure the protection of a given asset needs to be fully understood by all participants, and a person responsible for each behavior be assigned to every asset.

In order to ensure that, the critical and noncritical asset control behaviors must be fully and completely understood along with the required outcomes. Structurally, the control and organizational dependencies must all be identified and mapped, so that the people who execute the process can understand the various interrelationships of control elements and their interrelationships to the other parts of the solution as a whole.

This mapping will also include a description of each anticipated behavior for each baseline element as well as a value-based justification for their execution. This explanation is important for motivational purposes along with a listing of the controls that have been assigned and the justification for their assignment. The description should also unambiguously specify and document the constituent tasks. In order to ensure proper execution, these are the control tasks that must be performed by each individual and organizational element within that particular process. This description also aids in making resource allocation decisions.

Where standards have been utilized, the requisite practices, conventions, and metrics must be explained and justified along with how compliance will be monitored. The monitoring will include a specific specification of the testing standards and practices as well as the metrics that will be utilized to judge performance. Along with that, the individual tests and their methodologies as well as their metrics will be detailed.

The relevant reviews and audits will also be assigned to the verification of both the implementation process and the operational control infrastructure.

This assignment will stipulate the minimum number of reviews to conduct the audits that will be carried out and when they will be performed.

Finally, the method for problem reporting and corrective action must be itemized. That includes how any identified problems will be reported, and how nonconformances will be tracked and resolved, as well as who will be responsible for those roles. Over a period of time, the organization fine-tunes the controls that it believes are the most effective. Alterations are based on feedback from the stakeholders, and normally this involves a significant period of development time in order to make the process continuously improving.

Operational Measurement

The final concept is operational measurement. In a technical setting, this might be called a metrics program. At its core, the control infrastructure is created to implement a top-level set of assurance requirements. These requirements are derived from the security policies that are generally itemized in the cyber resilience plan and expressed as a formal specification of security requirements. The stakeholder perspectives, such as those of users and managers, are captured in those requirements and are characterized by a set of metrics.

These metrics express the abstract elements of the process in concrete terms. They are essential to ensure common understanding among all participants. For instance, a term like "threats" may have different meanings for managers and technical people. A network security person might characterize a threat as malicious code, intrusions, network interruptions, and denial of service, and would measure it by the things that affect them, such as instances, downtime, or mean-time-to-failure. Managers on the other hand might characterize threats in business terms, such as lost production, cost, and operational interruption time—they would measure it in the terms that are meaningful to them, dollar value, or cost.

Each is an appropriate point of view, and each perspective has a set of measures. The aim is to integrate these differing views and measurements into a single, meaningful, global understanding that meets the security goals. Organizational environment is a critical factor in this process because the rigor and application of the measures selected will vary. For instance, a highly secure government facility requires extensive and rigorous technical security metrics, while a large private organization might be more focused on the measures related to productivity and performance.

Measurement is a fundamental requirement of the cyber resilience control infrastructure implementation and sustainment process. That is because it is difficult to substantively impose any form of managerial control over a virtual asset. Thus, formal measurement programs provide the necessary data that help decision makers evaluate the control infrastructure's ongoing performance as well as maintain the necessary accountability.

As a result, every cyber resilience control infrastructure plan must be accompanied by a comprehensive description of the measurement program. Essentially, the measurement program ensures that confidence in the cyber resilience control infrastructure is maintained through objective data. Thus, the measurement process for any given cyber resilience control infrastructure must be able to provide consistent, data-based monitoring as a means of confirming that specified controls are in place and functioning properly. This is accomplished by performing regular reviews of operational elements at preplanned and mutually agreed on points in time.

Measurement programs are typically founded on a range of standardized metrics. These allow the organization to track and evaluate control performance for any asset in a given situation. Properly set up and maintained, the measurement program will provide critical tracking of the cyber resilience control infrastructure's overall operation and bring any undesired deviations to management's attention. Thus, an explicit plan for conducting validation and verification activities is an essential component of any control infrastructure implementation plan.

There are no universally recognized standard metrics for assessing the performance of security controls. Instead, individual organizations choose the measures that they feel best fit their particular situations. The rule for this is straightforward, whatever metrics are selected, they must be uniformly and consistently applied. In particular, since cyber resilience control operations are oriented toward harm to asset value, there is a need for a uniform definition and of what constitutes a loss of value. This requires the organization to delineate measurable characteristics that it considers to be indicative of harm or loss of value. That process has to be repeated for every potential security concern in order to make the data, produced by the assessment, accurate and meaningful.

The clarification of what constitutes value loss is especially necessary in order to help the measurement program function accurately. The organization can engineer or at least think through its metric requirements by clarifying the potential situations where unsustainable loss might occur. Consequently, it is possible to achieve a consistent and measurement-based description of the control actions using that approach. During the process of thinking this question through, the individual metric items can be identified that are appropriate to any set of assumptions about harm.

Maintaining the Cyber Resilience Control System over Time

The aim of a cyber resilience system is to ensure a dynamic and highly effective response to threats over time. To ensure the ongoing fulfillment of this purpose, the standard cyber resilience control processes must be continuously overseen, maintained, and reformulated if they are found to be deficient. The overall aim is to ensure effective oversight and understanding, reliable long-term sustainment, and

disciplined execution of the cyber resilience process. That status has to be maintained in the face of the challenges that arise out of the dynamic environment of organizational threat.

What we have discussed so far are infrastructure development concerns. The organization's cyber resilience architecture is the composite of all the controls that the organization has devised to provide active response and tangible prevention of threats. That control architecture must be maintained in a complete and trustworthy state. That includes the planning and installation processes that implement and assure the operational status of a formally designed and planned cyber resilience architecture. It also establishes and maintains the intentional interrelationships between the elements of the control set.

Nevertheless, standard policies, processes, and methodologies must define and be documented to ensure that the resultant cyber resilience process will continue to be appropriate and effective for a given threat environment. These policies and processes are altered as the threat picture evolves in order to maintain an effective cyber resilience architecture in alignment with the changes in the contextual situation. These policies, processes, and methodologies constitute the tangible elements of the cyber resilience architecture.

Besides creating a tangible architectural response, these policies, processes, and procedures also ensure continuous strategic improvement of the cyber resilience function as the threat picture evolves over time. It should go without saying that the cyber resilience architecture must be continuously maintained in alignment with the threat environment as changes. That involves the development, deployment, and continuous maintenance of the most appropriate and capable set of cyber resilience controls, interrelationships, and technical components.

Thus, the final stage in the cyber resilience process is to devise and implement a set of procedures that will assure the continuous assessment, creation, integration, and optimization of the cyber resilience control architecture. That architecture must be regularly and systematically evolved in order to ensure its currency and long-term effectiveness. To accomplish this, regularly scheduled organization-wide risk assessments must be carried out to ensure the continuous alignment of the cyber resilience architecture with the threat environment, and that architecture must be evaluated to ensure that it remains effective.

Since the field changes constantly, cyber resilience trends should also be assessed in terms of the way they might impact the evolution of the cyber resilience architecture of the organization. Potential control processes should also be evaluated in order to update and refine the cyber resilience architectural strategy. That might include the input from outside sources, such as consultants, in order to ensure that the broadest range of implications and requirements of the cyber resilience architecture are factored into the evolution. The aim is to ensure maximum awareness of the evolving threat environment.

New strategic directions also need to be evaluated as they relate to the ongoing development of the cyber resilience architecture and its substantive controls. From

this evaluation, rational decisions can be made about the most effective enterprise-level cyber resilience policies, processes, and methodologies. The aim is to maintain a dynamically effective cyber resilience and controls architecture in the face of inevitable change.

Logically, persistent, long-term control infrastructure sustainment is the mechanism that must be used to control change to the architecture of the cyber resilience control function. Formal control infrastructure sustainment provides two primary advantages. First, it maintains the integrity of all the elements of the control architecture. Second, it allows for the rational evaluation and performance of change to that architecture, as required. A formal sustainment process also gives the organization's decision makers direct input into the evolution of the protection scheme as it changes over time.

Once established, the control infrastructure sustainment process is maintained as an everyday operational process. The goal of the long-term sustainment process is to maintain the infrastructure control set at a defined level of correctness. Of course, that starts from the assumption that a complete and correct control set already exists. So initially, the documentation of the control set must unambiguously demonstrate that the current set of infrastructure controls is both trustworthy and also achieves the stated organizational purpose.

Typically, infrastructure sustainment underwrites the control set's ability to ensure confidence in the continued proper functioning of the cyber resilience function. Sustainment monitors the control set's ability to accurately identify and record problems, analyze those problems, take the appropriate corrective, adaptive, perfective, or preventive action, and confirm the capability of the system to ensure continuing assurance.

Control infrastructure sustainment does this by rationally controlling all changes to the form of the protection scheme. The management level authorized to approve those changes is explicitly defined as part of the overall process of cyber resilience planning. Changes at any level in the basic control structure must be maintained at all levels.

To ensure this, the most current cyber resilience methodologies, processes, and associated documentation must be identified and maintained in a stable state of assurance. Cyber resilience metrics must be developed and collected to support that assurance. These metrics should be used to improve methodology and process efficiency usage. The results of these analyses must be defined and analyzed. Cyber resilience metrics must be standard. They must be used for causal analysis to optimize the ongoing cyber resilience process.

Altogether, to minimize risks, the cyber resilience process needs careful attention to its personnel aspects. Formal teams must be established and coached in how to apply the organizationally standard methodologies and processes. Cyber resilience and control awareness, knowledge of policies, procedures, tools, and standards must be championed and promoted. Therefore, where necessary, cross-organization awareness, training and education processes must be established to

communicate the best cyber resilience practices to the employees of the organization. These materials must be appropriate and standard. Their coordination and deployment must be ensured by a formal process.

Chapter Summary

This chapter presents the ideas that underwrite the development of stable, large-scale organizational infrastructures for cyber resilience. It concentrates on the development of a conventional business process that will enable the creation and coordination of a stable, complete, and sustainable cyber resilience response. The goal of such a formal structure is to foster and maintain a complete and consistent environment for the conduct of the cyber resilience process. In order to create a proper infrastructure, the processes and practices that have been developed to fulfill a particular organizational purpose are systematized in such a way that they ensure the long-term, effective implementation of the overall business goal, which in this case is cyber resilience. The cyber-resilient response has to become a part of the overall day-to-day operation of the business in order to have any real impact. That is because every form of effective protection has to be continuously functioning in the everyday business environment.

Thus, cyber resilience must be embedded in the conventional business as an operational process. Most of today's corporate work is carried out in distributed, multilayered, multi-vendor, and even multicultural environments. All that work must be satisfactorily coordinated and controlled in order for the company's stakeholders to trust the outcomes. That coordination and control are often very difficult to achieve given today's complicated business challenges.

Given the inherent complexity in control formulation, the only way to ensure a systematic architectural solution is through a single coherent development strategy. The strategy must ensure a trustworthy, long-term cyber resilience capability, which will address all likely threats, and ensure that all items of critical and noncritical importance are secured appropriate to their potential value. A solution such as this must be systematically designed to merge all requisite behavioral and technical controls into a single, fully coordinated operational process across the organization. Security governance consciously builds an intentional structure of rational elements and interorganizational relationships that will sufficiently ensure the business' information assets.

A security architectural approach centers on the creation of a comprehensive organizational control system and security culture rather than building separate individualized security solutions for each problem. Essentially, the business defines a coherent framework that embodies all the necessary roles, their associated behaviors and practices, and the strategic policies necessary to achieve a single coherent, organization-wide cyber-resilient solution. There are five processes that underwrite the development of a security architectural framework.

The first of these principles is *scope*. In essence, scope addresses protection issues in the light of resource constraints. In the worst case, that might involve leaving a nonessential or low risk resource outside the boundaries of the protection scheme. Knowledge is obtained by formal assessment.

Assessment is the second essential element of the cyber resilience infrastructure development process. That is because assessment is a primary player in any form of risk-based security system development. The assessment function ensures that the organization fully understands the risks inherent in its security environment and the associated set of requirements.

The third element is *best practice*. Best practice is normally embodied in the form of a requisite best-practice reference model, which documents all the necessary controls of a standard governance model. In that sense, the detailed recommendations of a governance model comprise the basis for risk assessment.

Then there is *control definition*. Any one of these frameworks can be used to judge whether a security infrastructure and its constituent activities are complete and correct. Whatever the source, the reference framework should specify and populate a well-defined and detailed organizational best-practice architecture that allows the organization to develop its own explicit control policies, practices, and procedures.

The actual control system is *tailored* top-down to any desired level of detail for any given organizational application. A standard tailoring approach is particularly important in the creation of security infrastructures, since the implementation of the actual security system is always different in its particulars.

In order to be at their most effective, security infrastructure development processes should have both maximum flexibility and standard structure, which might seem like a contradiction. The simple resolution is to create the architecture of the solution top-down from the highest possible level of concept. The largest and most comprehensive view of the solution can then be used as a general classification structure within which effective controls for each of the specific security areas can be addressed. Using this approach, a specific cyber-resilient control architecture can be constructed for any given project at any given level of definition inside the model best-practice framework.

The outcome of the specific tailoring of controls to satisfy each of the general security categories of the standard should be an operational process model, that embodies best practice, and which specifically ensures cyber resilience for a given organizational application.

The behavior of any form of control system has to be documented. Documentation is an absolute ongoing requirement. In many respects, the documentation of control outputs is the absolute bedrock on which oversight and management rest. So, it would not be too much of a stretch to conclude that it is the documentation process that ensures that the cyber resilience function is properly run.

Practical infrastructure design involves identifying and documenting all the requisite elements of the eventual solution and then specifying how they

interrelate. Any technical, human, or practice behavior that does not achieve the desired organizational purpose would be a waste of money. Therefore, the process of developing the right mix of elements requires the organization to strike a rational balance between the outcomes of the proposed infrastructure component and its associated cost.

That process is architectural in the sense that actual structures of related elements are being assembled based on some inherent logic. The activities that take place within a cyber resilience system embody a complex network of input-output relationships, which taken together comprise the single process-based cyber resilience assurance solution. To be functionally correct, the architectural model must specify an exact set of exit conditions. That allows the organization to determine whether the specified task has been executed properly in that instance.

The standard architectural model that provides the basis for developing the applied system must encompass and describe the entire structure of the solution, from top to bottom. Therefore, it is necessary to define every process and practice at a level of detail sufficient to ensure proper operation of the cyber resilience solution. All the details of the elements that are used to categorize that solution must be expressly traceable and derivable from the higher-level elements of the architecture.

The basic unit of an architectural process model is the task cell. In practical security implementations, the task cell is also known as a "control." Task cells are unitary functions. That is, each task cell is specified as the means to carry out one specific task, and one task only. The actions the cell defines have logical entry conditions that are required for proper task initiation. These include the inputs from any prior activities at all levels of abstraction. They also must produce an intended outcome from those inputs.

Once a complete set of standard process cells has been defined, a process model can be constructed. This is done by interconnecting the basic set of task cells in various ways to produce a tailored set of best-practice controls that satisfy the needs of a given application. The practical mechanism for ensuring this requires the following five steps:

First: A standard model of best practice must be adopted to define the general form of the architecture, and for tailoring out a comprehensive solution.

Second: Once a standard model has been adopted, a set of task cells is tailored out of the general activity recommendations.

Third: The specific ETX criteria and expectations are specified for each cell.

Fourth: The precise specification of the required input and anticipated output behaviors is done. This allows the organization to monitor and track the behavior of each cell.

Fifth: The requisite monitoring is described in an assessment plan. The standard and systematic assessment, measurement, and reporting that are described in that plan are then carried out on a systematic operational basis.

There must be a systematic means to ensure the monitoring and control of the evolution of the cyber resilience process over time. The ongoing understanding of the explicit operational behavior of the formal array of cyber resilience controls is derived from a continuous monitoring process. That process will allow the organization to benchmark control set performance over time. Ideally, management will be able to continuously evaluate the performance of its cyber resilience control infrastructure, and then take substantive action to ensure its effective operation in real time.

The practical assessment is performed in precise measurement driven terms. Depending on the outcome of that assessment, there might be a long period of trade-offs and refinement before an eventual decision can be reached about the effectiveness of the cyber resilience function. However, the final documented solution must objectively demonstrate that it addresses the strategic goals and business objectives of the organization.

The aim of a cyber resilience system is to ensure a dynamic and highly effective response to threats over time. To ensure the ongoing fulfillment of this purpose, the standard cyber resilience control processes must be continuously overseen, maintained, and reformulated if they are found to be deficient. The overall aim is to ensure effective oversight and understanding, reliable long-term sustainment, and disciplined execution of the cyber resilience process. That status has to be maintained in the face of the challenges that arise out of the dynamic environment of organizational threat.

Once established, the control infrastructure sustainment process is maintained as an everyday operational process. The goal of the long-term sustainment process is to maintain the infrastructure control set at a defined level of correctness. Of course, that starts from the assumption that a complete and correct control set already exists. So initially, the documentation of the control set must unambiguously demonstrate that the current set of infrastructure controls is both trustworthy and also achieves the stated organizational purpose.

Keywords

Acceptable risk: A situation where either the likelihood or impact of an occurrence can be justified

Analysis: An explicit examination to determine the state of a given entity or requirement

Architecture: A designed entity with a consciously designated purpose

Assurance: The set of formal processes utilized to ensure confidence in software and systems

Baseline: A collection of entities related by a similar purpose and time frame

Baseline security: A minimum level of acceptable assurance of proper performance

Benchmarking: Measurement by comparison with standard measures or past performance

Boundaries: A perimeter that incorporates all items that will be secured
Business impact analysis: An assessment of the effect of the occurrence of a given event
Common features: Functionality shared among a number of standard practices
Controls: Activities built into a process designed to ensure a particular purpose
Infrastructure management: Role that oversees and maintains a defined process architecture
Likelihood determination: An assessment of the probability that an event will occur
Monitoring: Specific oversight created by a planned collection and analysis of data
Quantitative management: Decision-making that is supported by empirically derived data
Prioritization: Assignment of importance based on the perceived value
Process: A collection of practices designed to achieve an explicit purpose
Process architecture: The long-term method of organization of the overall information and communication technology (ICT) work
Process specifications: The explicit work rules and requirements of a given operation
Risk: A given threat with a known likelihood and impact
Risk behavior: An action that will produce a defined exposure or risk for the organization
Risk level: The likelihood and impact that are considered acceptable before a response is required
Risk treatment: The specific control response that is planned for a given adverse event
Security controls: Mechanisms designed to ensure proper performance of the process
Security metrics: Quantitative measures of security performance
Security system: Formal collection of controls aimed at mitigating all known threats
Security testing and evaluation: Validation of the secure performance of the control process
Security vulnerabilities: Explicit known weakness that can be exploited by a given threat
Specification: Documentation of explicit requirements of a given system
Stakeholder: Person responsible for a given item or function, generally also the decision maker
Strategic alignment: Assurance that the security actions of the organization directly support their goals
Strategic framework: The generic organizing and control principles that an organization uses to underwrite the management of its information function
Strategic management plan: The prescribed activities to achieve the long-range intentions of the organization

Strategic planning process: A set of rational activities that are undertaken to ensure long-range directions of the organization

System design: Assurance that the architecture of the system meets requisite criteria

Technical (functions): Automated mechanisms designed to ensure secure performance of the process

Testing: Validation of performance of a piece of software or system

Threat: Adversarial action that could produce harm or an undesirable outcome

Vulnerability: A recognized weakness that can be exploited by an identified threat