Chapter 1

Introduction to Organizational Security Risk Management

At the conclusion of this chapter, the reader will understand:

- The role and importance of risk management in the cybersecurity process
- The issues associated with risk and generic risk management
- The form and content of the risk management process
- The general structure and intent of risk-oriented frameworks
- The general application and development of a risk-based strategy
- The generic elements of the risk management process

1.1 Introduction to the Book

The goal of this book is to provide a comprehensive understanding of the strategic risk management process as well as the underlying principles and a standard risk management framework. Risk management entails a formal set of steps that are carried out to protect an organization's assets from harm that may be caused by inadvertent or deliberate acts of destruction. Risk management involves a systematic architecture comprising all the necessary controls to prevent unauthorized use, loss, damage, disclosure, or modification of organizational information. Specifically, this chapter discusses the formal processes for identifying, managing, and mitigating risk as prescribed by the National Institute of Standards and Technology's (NIST) risk

management framework (RMF). In this chapter, we also discuss the general uses for the framework and the contexts in which it applies.

In some respects, this book is as much about standardization as it is about risk management. Hence, Chapters 2 and 3 present an overview of the role of the standardization process in ensuring a consistent response to a given issue of importance. This includes a discussion of why information assets are difficult to protect as well as the part in which commonly acknowledged best practices apply in ensuring an informed response. The discussion will also center on how to use the NIST's RMF as a standard means of deploying an appropriate set of information technology security controls. We lay out the issues involved in implementing a standard process, including the benefits that derive from it, as well as potential pitfalls. We also try to give you an understanding of the implementation process, which is best demonstrated by applying the RMF to a specific context.

1.2 Risk Is Inevitable

Risk is a fundamental element of human life in the sense that risk is always a factor in any situation where the outcome is not precisely known (Figure 1.1). In addition, the necessary calculations that we make about the probability of some form of harm resulting from an action that we take are generally a given in our decision processes. Whether the risk assessment involves decisions about a major corporate

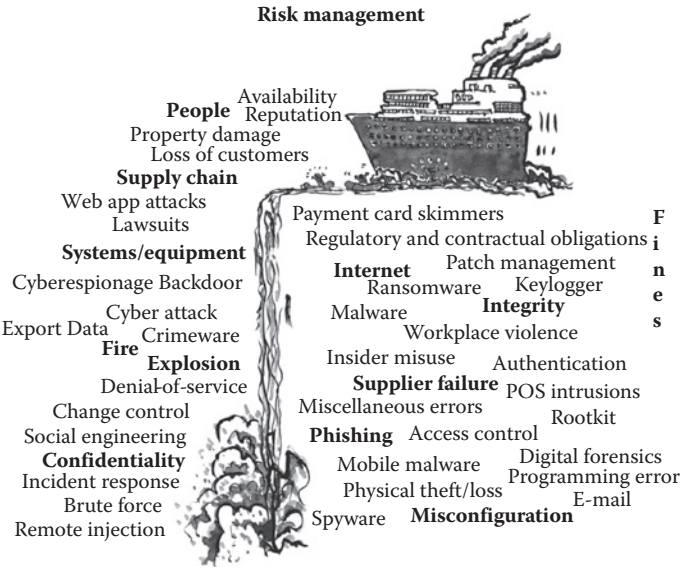

Figure 1.1 Security risk management.

initiative or just making the decision to walk down the street, we are always anticipating, identifying, and evaluating the potential risks involved. In that respect, we can be said to be constantly managing risk in everything we do.

The reason why risk management is a particularly important aspect of the cybersecurity body of knowledge (BOK) is that information and communication technology (ICT) and information assets are more difficult to account for and control than most conventional physical assets, because ICT involves the production and management of virtual, highly dynamic products, which makes it difficult to identify what to secure, let alone how to do it. That puts risk management center stage in the consideration of how to establish and maintain a secure ICT environment.

By definition, ICT assets are something of value to the business. The risk management process specifically ensures the assurance of three generic protection criteria, as shown in Figure 1.2. These three criteria assure against meaningful loss of *confidentiality*, loss of *integrity*, and loss of *availability* (CIA).

From a security standpoint, the most logical generic criterion might be assurance against a loss of confidentiality. *Confidentiality* is a security principle that encompasses an organization's requirement to restrict access to any sensitive information or data that it keeps. Obviously, if the organization's data and information could be made public without risk, there would not be a need for this attribute; however, this is rarely the case.

From an operational point of view, confidentiality is founded on establishing and adequately enforcing access control. Data and information are essential to the business operation. And in many information-intensive organizations, it might be the only real asset that is kept. For instance, most financial data within a company is sensitive and

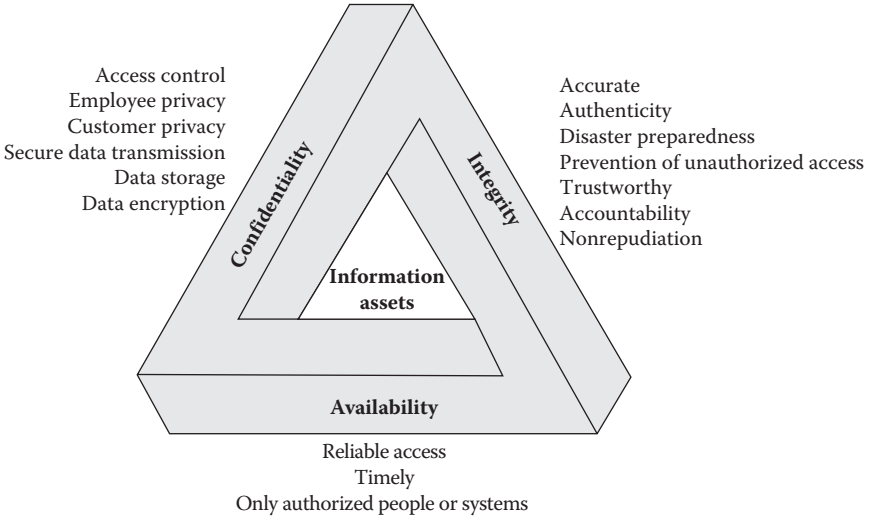

Figure 1.2 The confidentiality, integrity, and availability (CIA) triad.

access is almost always rigorously safeguarded. So, one way to view the monetary value of confidentiality is to imagine how much competitors might pay to have access to the data and information of a company or the cost of litigation if a legal requirement was violated. Thus, in that respect, the organization has a legal and ethical requirement to protect its sensitive business information as well as employee and customer privacy.

The second characteristic is *integrity*. The integrity of data or its attended processes is determined based on how authentic, accurate, and complete the data is. It is easy to appreciate the value of integrity in the context of financial business transactions. For example, if a bank could not depend on its account balances, it could potentially sustain a large loss by disbursing checks not covered by actual funds. In an inventory system, there is the potential to lose expensive materials if the counts were inaccurate due to faulty data. Or publically, the release of unreliable data that is used as background for a damaging story might expose a newspaper to legal action.

The third characteristic, *availability*, ensures that information is provided to an authorized user when it is required. The best way to understand the value of availability is to ask, "What would happen if the information was not available to support a given action or decision?" For example, what would happen if the business' payroll data were erased on payday? If the payroll program were suddenly inoperative, no one in the organization would be paid as expected. Imagine the chaos in a company the size of General Motors or IBM if they were unable to pay their employees or suppliers when they needed to. Given the potential harm that each of these principles might represent, all of the meaningful risks in each of these areas must be rationally managed.

Because every organization is unique and implements security differently, the actual process to identify, evaluate, and ensure that the meaningful risks in each of the CIA areas are properly managed generally involves the same eight requirements, which are as follows (Figure 1.3):

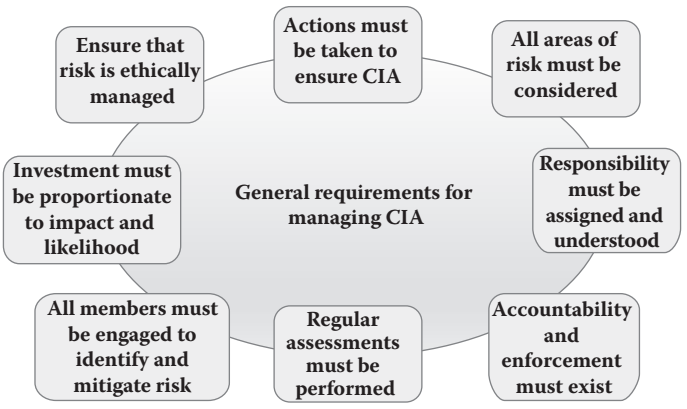

Figure 1.3 General requirements for meaningfully managing CIA.

1. Identifiable actions must be taken to ensure correct, confidential, and available information.
2. All relevant areas of risk must be considered in any given solution.
3. The responsibility for risk management must be explicitly assigned to individuals and understood.
4. A system of accountability and enforcement for risk control must exist and be documented.
5. Regular and systematic assessments of risk status must be performed.
6. All members of the organization must understand the importance of and work to identify and mitigate risk.
7. Investment in risk management must be kept proportionate to the impact and likelihood of the risk occurrence.
8. The organization must ensure that risk is ethically managed.

In practice, organizations should design, implement, and follow a systematic process to establish a persistent operational risk management process. This design and management process is a strategic activity in that it involves short- and long-range considerations. Thus, planning for strategic risk management is necessary in order to ensure continuous risk assurance. And a formal strategic planning process is necessary to implement an organization-wide risk management process. Risk management itself must incorporate all of the elements of the business within its scope and the process should reach to the boundaries of the organization.

The outcome of the implementation of a risk management process is a concrete organization-wide risk management scheme that is documented. The risk management scheme will balance the aims of a long-term risk control policy with real-world conditions and constraints. The atomic-level components of the risk management process are a set of substantive security controls that ensure the requisite level of assurance against loss. These security controls should be traceable directly to the individual policies that defined their need. This is a closed-loop process in that the ongoing alignment of risk security controls to individual policies fine-tunes the evolution of the substantive risk management process and ensures its effectiveness in the operational setting.

One problem is that the term "risk management" is rather nebulous. So, the overall process itself requires a definition of what risk management means. A concise statement and commitment to the work is needed in order to make the practice standard. Standardization is important because a lack of effective, coordinated implementation and execution of the process has made overall risk management efforts ineffective. Worse yet, employees might feel the effort is the "flavor of the day" and not take it seriously. One does not need to look any further than the increasing number of incidents in cyberspace to confirm that.

The lack of coordinated action has been so pervasive that a logical response is the formulation of a comprehensive and coherent specification of the commonly accepted best practices for risk management. The specification could then be used

to guide the creation of an effective risk management scheme for all organizations. In that respect, steps were taken by the federal government to formally research and develop a standard and comprehensive risk management process.

The specification of commonly accepted standard processes is the role of the NIST, the U.S. government's standards making body. Of specific interest here, the NIST has developed and published a formal reference model for the management of risk simply called the RMF, as shown in Figure 1.4.

This large-scale standard model serves as both the specification of a fundamental process for understanding the risks involved in assuring information and ICT organizations and the foundation for deploying the common control mechanisms required to manage the risks that exist within them. It has the additional advantage of providing the umbrella definition of the processes for achieving Federal Information Security Management Act (FISMA) and NIST certifications.

An important justification for this standard is that the RMF also defines the basis for a comprehensive strategic governance approach to risk. A governance rather than a technical approach is a highly advantageous strategy because, notwithstanding the issue of whether the cybersecurity function itself can ever fully embrace all of the issues associated with assurance, a governance-based solution is more easily understood and acceptable to the managers and nontechnical people who comprise the majority of the organization.

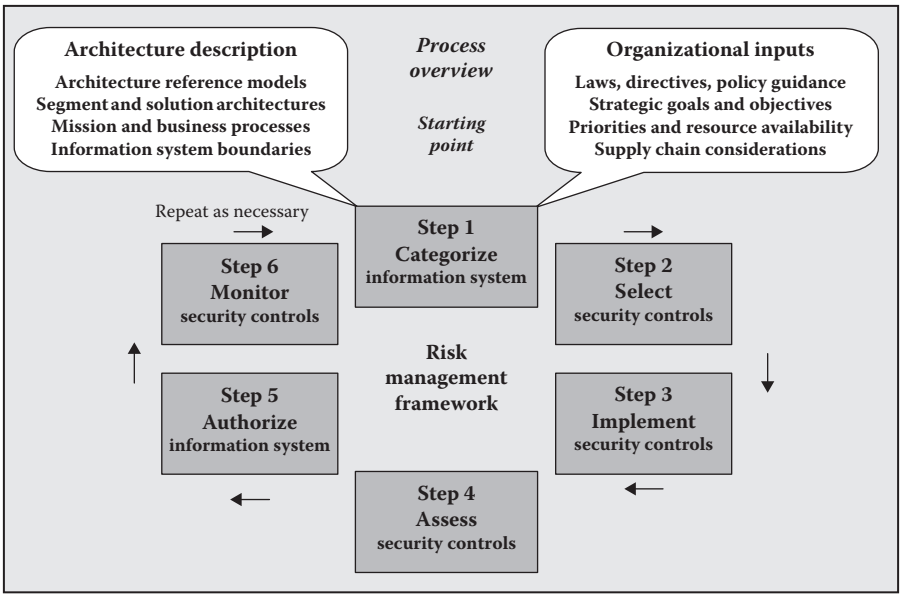

Figure 1.4 The National Institute of Standards and Technology's risk management process overview.

A governance approach is appropriate for any organizational setting. In essence, a generic governance model constitutes a flexible top-down organizational process for establishing persistent risk management actions and the formal selection and maintenance of appropriate security controls. Moreover, since the RMF is founded on an established policy and procedure approach, it is able to capture and communicate the nature of the specific risks that an organization may encounter. And finally, since the framework itself is built and maintained through a comprehensive identification and assessment process, it can assist in rationally and systematically identifying changes in the threat environment as they occur.

1.3 Strategic Governance and Risk Management

Starting from the assumption that a standardized risk management process should be applied organization-wide (which is what we believe), risk management is a strategic issue, rather than a narrow technical concern. The reason to adopt an organization-wide risk management approach is to avoid the dysfunctional effects of a typical piecemeal solution where every department is managed by its own commonly accepted business practices. These are often based on an individual unit or manager's ideas about the proper way to accomplish a particular organizational goal. And regardless of whether they are universally standard or documented, these become the corporate way of doing business. One problem is that those approaches are often not coordinated effectively in the operational environment. In some cases, they can actually cause dysfunctional conflicts. And corporate risk management has often evolved this way. Organizations develop specific one-at-a-time responses as risks present themselves, rather than addressing them by employing a single, coordinated management strategy. Moreover, as new risks appear in the corporate threatscape that have not been seen before, they are not incorporated into any specific management techniques that the organization employs to mitigate and contain them.

The alternative approach to piecemeal risk management is a formally defined and instantiated architecture of comprehensive risk management best practices, which are specifically aimed at optimizing risk controls within the company. As with any complex system, formal risk management practice can only be implemented through a rational and explicit planning process. The planning activity fits the strategic purposes and responsibilities of standards-based risk management to the security needs of the organization. From the standpoint of the rest of this text, it is the creation of that strategic risk management capability, which the RMF leverages, that will drive the presentation and discussion of the framework.

Risk management is basically built around information. In effect, risk management gathers and utilizes information from all sources, in order to decrease the possibility of future risks. The information-gathering activity is aided by a set of formal processes and technologies. And, at its core a successful risk management

function relies on the ability to assure that the processes, practices, knowledge, and skills of risk management are incorporated as quickly and efficiently as possible into the organization's substantive decision-making processes.

In addition to providing the information that helps guide strategic decision-making about risks, the risk management process also makes certain that a commonly accepted and systematic set of policies and procedures are in place to handle known risks. That responsibility is operationalized through a standard set of operating procedures. Those procedures ensure that the risk planning, analysis, response, and process management function are always directly aligned to the goals of the business operation. Nevertheless, the primary purpose of risk management is to ensure a disciplined and systematic response to the risks that the organization considers a priority.

1.4 Elements of Risk Management

In simple terms, the risk management process assesses the likelihood that any given action will adversely impact something of value to any given entity. That includes things of personal value such as money, health, or even life. Once those risks are known, the risk management process deploys all of the measures that are necessary to ensure that consequent harm does not occur.

Some organizations manage risk in a highly quantified and data-driven way, for example, corporations that require high levels of integrity in their products as well as the segments of the critical infrastructure where the potential failure of a crucial system could result in a set of highly unwelcome consequences. Others tend to spend less on risk management and spending levels are influenced by the nature of the threat environment and the value and sensitivity of the assets that are being protected.

Because identification and understanding are such important aspects of risk management, assessment provides the fundamental focus of the process. Risk management is operationalized by a continuous process of assessing the organizational environment aimed at identifying and understanding all of the potential threats and the negative impacts that might affect the business. Once these have been identified and characterized, specific steps are then devised and implemented to mitigate any adverse outcomes.

Given its focus on the support of substantive decision-making, an important underlying factor in risk evaluation is the uncertainty principle. Uncertainty is a key element in assessing threats because risk entails future consequences. In essence, the outcomes of any given threat have to be fully understood in order for an intelligent decision to be made about the way forward in addressing it. However, there are usually a number of unknown, and therefore unevaluated, factors that might be associated with a given threat. Thus, the institution of standard and persistent identification, understanding, and response practices becomes an important element in the risk management process.

It goes without saying that it is easier to identify and evaluate risk in less complex environments. Yet, every aspect of cyberspace is abstract and complex. Therefore, risk management for cybersecurity requires a much different approach to the understanding and evaluation of risk. The process in the virtual world has to touch on factors than would normally not be part of the decision-making processes in the conventional physical world—such as how to authorize the acceptance of an invisible product. Accordingly, the sheer virtuality of ICT environments alone poses a threat.

The issue of threat management is important to our existence as a nation because ICT is the platform on which our modern society rests. Consequently, the huge increase in the number of strategic threats to computers and networks is a compelling danger to our modern way of life. The generic areas of threat have been variously categorized into terms such as "cyber-crime," cyber-terrorism, and "cyber-war." And in response to all of this turmoil, the past 15 years have witnessed the creation and evolution of a specialized new profession that is dedicated to addressing the many novel risks of the virtual world. The aim of that profession is to assure that ICT systems and the information that they contain, process, and communicate are protected against all logical forms of unauthorized access, use, disclosure, disruption, modification, perusal, inspection, recording, or destruction. That profession is presently termed "cybersecurity."

Cybersecurity evolved out of the practices and procedures of the older discipline of information assurance. One aspect of the original discipline was the responsibility to manage all risks related to the use, processing, storage, and transmission of information or data and the systems and processes used for those purposes. Cybersecurity incorporates a holistic approach to protection in that all aspects of risk mitigation in virtual and physical space have to be included in the protection scheme. This includes the creation and deployment of a complete and appropriate set of electronic and behavioral countermeasures.

This requirement is not simply a computer science challenge. It requires knowledge and practices from a wide range of traditional security fields, such as continuity management, forensics, audit, management science, software, and systems engineering, and even fields such as law and criminology. Consequently, what is required to manage cybersecurity risks is a complete and provably effective framework that ensures the proper coordination and use of all appropriate methods in the execution of the process. The framework should be expected to consolidate provably correct approaches into a single logical and coherent model of operation. The model contains all of the commonly accepted security best practices necessary to provide effective mitigation and management of all known risks to individuals, operations, and assets of the organization.

The key concept is "commonly accepted." A commonly accepted model of best practice establishes a standard point of reference. A unified vision is necessary to establish coordinated actions in the management of risk. Comprehensive coordination is a necessity because *all* potential risks must be identified, assessed, and

responded to at all levels of the organization. The necessity for a complete, unassailable solution is a problem for the average manager. That is because conventional managers simply do not have the background or training to identify every potential risk, let alone devise foolproof methods to mitigate them. Nevertheless, particularly given the level of skill and sophistication of the large collection of malicious agents out there, it is critically important to implement comprehensive organization-wide protection since any system with an exploitable hole is a potential hazard.

As a result, there has always been an implicit requirement for the profession to establish and maintain a standard and comprehensive point of reference that practitioners can utilize to structure a practical risk management solution for their specific situation. Consequently, it is an attractive idea to consider employing a single commonly recognized standard, which specifies a single effective method for risk management.

Nonetheless, another underlying issue is how to get the most effective assurance out of the organization's limited resources. Any risk can be managed if enough money is thrown at it. However, no organization has the wherewithal to effectively put a cop on every street corner, so to speak. So, managers must weigh and balance the deployment of their risk response against the potential likelihood and material consequences of the threat. In day-to-day commercial operations, this means that it must be possible to make an informed decision about the level of risk that can be acceptable for every given situation. And given its layers of complexity, this is a particularly difficult task with cybersecurity risk, especially when the decision is weighted against the possible cost of failure.

Consequently, a coherent set of best practice methods, which let decision-makers benchmark existing and planned risk management resource usage, using the most expert advice available, is an important strategic management tool. This is because the drive for competitive advantage and the need for cost efficiency have driven corporations toward a growing dependence on technology. And thus the impacts of ICT risks have become an increasingly critical factor. Moreover, given that technology experiences rapid and dynamic change, the BOK regarding risk management must be deliberately researched, publicized, updated, and maintained. That condition justifies the role of the NIST in the development and promulgation of guidance about risk management.

The NIST's RMF was designed to offer a structured, yet flexible, means for analyzing and deciding how to alleviate the risks that arise from the information systems within an organization. The idea of adopting a coordinated set of formal risk management practices is a relatively new concept. Cybersecurity risk encompasses all of the risks that relate to the use of ICT. Thus, the risk management approaches that are specified in the RMF are intentionally broad-based. This is because those recommendations are meant to dictate how to assess risk and employ the appropriate risk mitigation strategies for all conventional ICT organizations.

This requirement implies the need for a single umbrella model that defines the elements and relationships of the risk management process. The specific steps

for risk management take place within the structure created by this overarching model. And these are captured in the appropriate supporting NIST and security standards and guidelines that apply to that particular problem. The framework was derived from and builds on the collection of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC), the Institute of Electrical and Electronics Engineers (IEEE), and NIST standards. It also consolidates information from various standard body publications, such as the Committee on National Security Systems Instruction (CNSSI) and the Department of Homeland Security Federal Continuity Directive 2 (FCD 2), and provides examples of ways to implement those standards and guidelines.

1.5 Risk Types and Risk Handling Strategies

There are four strategies that are generally employed in dealing with risk. The first strategy is to *accept* the risk and consequent losses. The second strategy is to *avoid* the loss by performing the necessary actions to eliminate the risk. A third strategy is to *mitigate* or *reduce the effect* of the risk. The last strategy involves *transferring* the risk to another party. That transfer can be achieved through contracts, insurance, or a variety of similar mechanisms. Nevertheless, no matter what approach is used, the organization has to adopt a formal strategy to decide how to address each of its risk categories. Likewise, regardless of the circumstances, the decision about what to do about the risk is purely in the domain of the designated decision-maker(s).

Accepting risk and the consequent losses is the most common approach for risks that rarely occur or where there is limited harm. Many risks pass unidentified or unacknowledged through the corporate risk management function because the cost of addressing the risk would not justify the potential cost of the harm. The decision to accept a risk can also change as the risk situation changes. After all inherent risks have been addressed by controls, there is still risk left over and an organization may decide to accept those risks. Even though the potential for harm exists, the present harm from the risk has been judged to be acceptable. Therefore, residual risks are still identified and tracked through the risk analysis process.

Risk avoidance is aimed at preventing the risk from actually occurring. Information security has three standard components: *prevention, detection,* and *response*. The prevention element and all it involves are examples of risk avoidance. Training programs, which are designed to increase the ability of employees to recognize and respond to incidents, are good examples of this type of risk handling approach. The information security process is heavily geared toward avoidance in order to reduce, as much as possible, the amount of harm by addressing the risk directly.

The last two components of the information security process, detection and response, are embodied in the risk mitigation and risk transference approaches. In the case of risk transference, the response requires an outside party to assume the impact of the risk. Insurance is a prime example of this type of assumption.

Obtaining insurance against specific risks does not prevent the risk from occurring, but it provides financial reimbursement to make up for a loss that will occur. Risk transfers work well when the risk is associated with a financial loss. Risk transfers are less effective when the loss is associated with less tangible things, such as customer service/retention, organization reputation, or in some cases regulatory requirements.

Risk mitigation approaches are the steps that an organization takes to minimize the potential loss in the event of the occurrence of a risk. For instance, an intrusion detection system will not prevent someone from actually intruding on the network. Instead, intrusion detection systems function as "burglar alarms" to limit the time that an intruder is allowed to roam undetected through a network. The limitation of time will not prevent damage. Instead, the limitation of time is meant to restrict the damage that might occur.

An important feature of the RMF is that it provides a practical basis for developing and maintaining comprehensive risk management controls for all aspects of a business's information assets. The objective of the RMF is to provide a common sense basis to develop, implement, and measure effective risk management practices. It is implemented through an organization-wide participative process and any business that has faced compliance issues with FISMA or NIST should be able to easily follow the RMF process.

The goal of the RMF initiative is to define and communicate a commonly accepted and standard basis for building risk management best practice. The RMF scheme compares the risk management practices of an organization against the threats and vulnerabilities it faces and prescribes a systematic mitigation approach for those threats (Figure 1.5).

It is designed to enable ICT managers to leverage their levels of risk awareness to a higher status. It allows companies to identify gaps in their risk management processes. It also allows companies of all sizes to demonstrate the effectiveness of their risk management program to prospective trading and investment partners. The RMF model underwrites assurance of risk management capability to any outside entity because it provides auditable and certifiable evidence that a scheme is in place to mitigate them.

Organizations have to document that they have considered the risk to their assets and have control measures in place to protect themselves against it. Those measures themselves are commonly understood as correct and specified in the NIST Special Publication (SP) 800-53 Revision 4 Standard, *Security and Privacy Controls for Federal Information Systems and Organizations,* which is the basis for verification of compliance to the FISMA. And in that respect, the RMF provides the risk-based assessment model for deploying the controls necessary to obtain formal certification of compliance with both FISMA and NIST. From a marketing perspective, certification to the RMF can also provide a basis for brand differentiation for ICT products. In that respect, the presence of an audited and certified security system becomes a true means of demonstrating the commitment of an organization to proper cybersecurity protection.

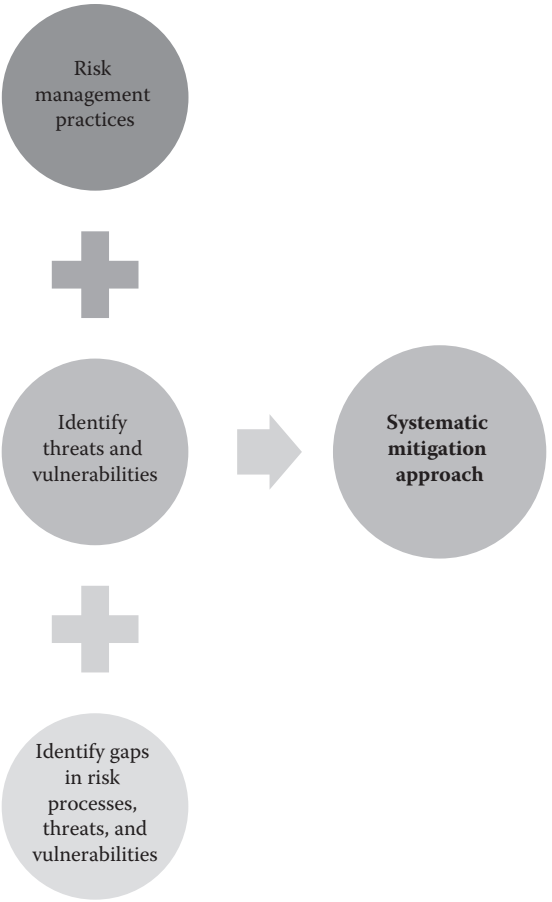

Figure 1.5 Risk management initiative goal.

Finally, in the cases where FISMA or other forms of audited proof of compliance are required, the external auditors will be able to determine that the organization has adopted a commonly accepted means to deploy a standard cybersecurity system within their organization. That is likely to make the certification and accreditation process a whole lot more efficient, as well as support the company's stance on any outcomes that could be called into question during the audit.

Besides these advantages, the RMF approach also offers some operational benefits. Cybersecurity tends to be tactical and reactive by nature, waiting for the bad guys to show up. On the other hand, if the defense-in-depth scheme is based on deterrent principles that are complete and comprehensive, the array of countermeasures can be protective rather than reactive. Organizations can initiate a full-scale set of procedures designed to prevent rather than remediate threats and work more proactively.

The RMF process supplies the management basis for identifying and organizing the comprehensive set of common best practices that the organization needs to establish and maintain control over its ICT risks. Since the RMF was designed to meet the needs of a range of target constituencies, and it is applicable to a range of ICT environments, it has the potential to deploy all of the necessary cybersecurity assurance elements to ensure an organization's systems are protected throughout their life cycle.

The RMF applies equally to building assurance as well as the long-term maintenance of assurance for information assets, embodied in organizational ICT systems. The activities in the RMF apply independently whether the actual system development and maintenance work is performed internally or externally to the organization—for example, outsourced. The risk evaluation approach applicable to the definition of a cybersecurity solution for a single system or multiple sites may even be applied on a shared basis between multiple parties. It delineates all of the elements of risk assessment that are necessary to structure a complete security response for any organization. This can be captured and expressed in everything from informal agreements up to a legally binding contract.

Since the RMF touches on every aspect of how to assess and manage risk, it forces companies through a step-by-step evaluation of their needs and responsibilities with respect to their ICT function. Nevertheless, the process itself is generic. That is, it provides the direction at the control level and not the step-by-step procedures necessary to manage risk. Thus, the generic assessment and implementation approach must be adapted to fit each given situation.

In essence, an optimum approach is engineered out of the RMF model for each individual organization. The understanding of risk that the RMF provides and the appropriate set of control objectives selected from NIST SP 800-53 Revision 4 comprise the actual form of the eventual response. Accordingly, the approach to implementing the RMF is hierarchical. Or in essence, an explicit cybersecurity solution that includes step-by-step policies and procedures is developed for each control area, at any level of definition top-down within the reference model provided by the RMF. And in that respect, the RMF assumes that specific cybersecurity approaches will be tailored to the outcomes of the common assessment process that is specified within the framework. This is accomplished in three steps. Once the threats, vulnerabilities, and weaknesses that the organization faces are assessed and their likelihood and impact are determined, policies are defined for each applicable control area. This serves as a foundation for tailoring.

Then, explicit control specifications are defined for each of the applicable areas of security risk management using the control recommendations of NIST SP 800-53 Revision 4. Finally, the real-world, day-to-day procedures/individual tasks are tailored and detailed for each individual role within the risk management process. These work instructions substantiate the standard behavioral specifications for a

particular area of identified risk. The end result is an explicit set of risk management actions, which are based on the standard but accommodate all known threats.

Substantively, the actual operational response requires precise identification of the organizational context and requirements associated with each risk. Then a control is tailored that addresses those contextual requirements in the most effective way possible. Because risk contexts normally impose singular behaviors, the control procedures are usually tailored and implemented at the project level in various, project-specific ways. However, the definition and overall control selection process is executed globally for the entire business. The idea in all of these cases is to build a practical solution that will address the known threat environment, while continuing to incorporate the best practice recommendations of the framework.

This hierarchical approach creates a tangible, complete, and rational architecture of cybersecurity controls. It is imposed top-down directly out of the threat space into a precise set of security policies that define the organization's overall risk response. That definition process then continues through the practical management activities that implement these policies, right down to the level of utilitarian tasks. Tailoring can then be finalized by identifying the unique risk management issues, problems, and criteria for each instance and then making the necessary execution adjustments to fit the overall risk strategy.

The outcome of this tailoring process is a set of explicit behaviors, which become the tangible instantiation of the cybersecurity risk management scheme within any given organization. In general, the tailored set of procedures is the most visible and useful to the line manager, because it makes the recommended standard operating procedure (SOP) concrete in day-to-day practice. Moreover, the tailored set of best practices embodies and conveys the exact substance of the assigned activities and tasks for personal risk management behavior to every one of the employees working within the organization as a whole.

In concept, the controls itemized in NIST SP 800-53 Revision 4 are the general basis for tailoring out explicit control behaviors. But these control recommendations are not stand-alone elements. They are actually one facet of the aggregate set of best practices, which when properly arrayed as a set of standard activities, produces a rationally managed risk function within any organization. The controls form a complete and tightly integrated system as a set; however, in order to fulfill any aim or purpose that it might have, organizations can choose an appropriate subset from the complete set of NIST SP 800-53 Revision 4 controls.

1.6 Overview of the Risk Management Process

The steps to establish a standard risk management process involve five generic organizational functions: identification, assessment, control selection and implementation, test and measure, and continuous monitoring, as shown in Figure 1.6.

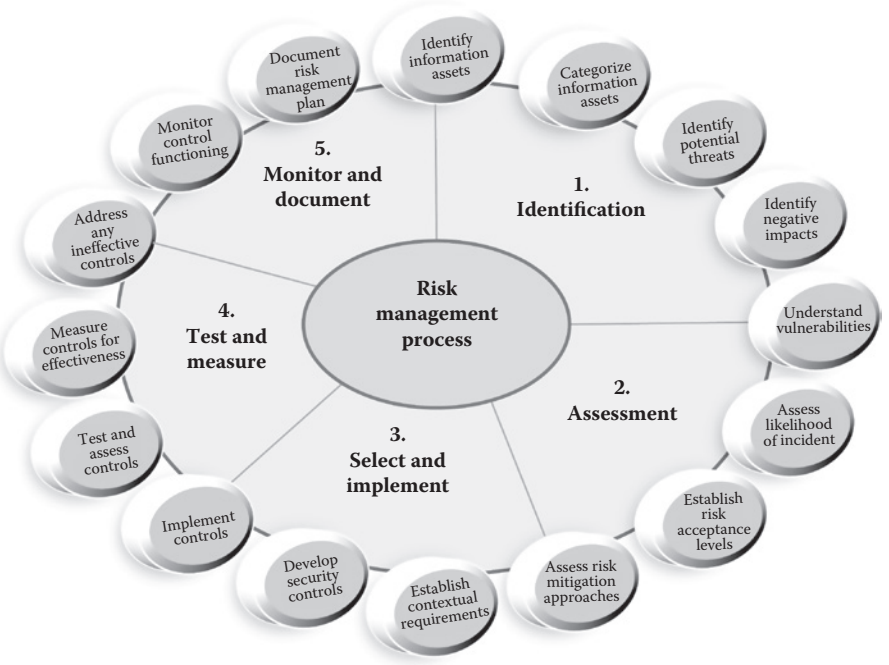

Figure 1.6 Risk management process.

1.6.1 Establishing the Risk Management Planning Process

The risk management plan shapes the risk management process. The primary role of the risk management plan is to create the framework for the detailed policies and procedures that will comprise the risk management process for the particular organization. The top-level risk management plan provides the strategic context that is needed to ensure that the organization's overall business objectives and goals are understood and then factored correctly into the decisions that are made about risk.

In that respect then, the overall plan for risk management needs to be crafted in broad, organizational terms, with the specific details of the approaches to be adopted left to lower-level operational plans. It is important that this high-level document defines the comprehensive processes and interrelationships needed to build a complete picture of the organizational risk situation. The ideal would be to create a roadmap that will let executive managers develop the strategies they will need to address existing risks.

First, the risk management plan should document the roles and responsibilities of the risk management team. The assignment of responsibility should be stated at a high enough level to allow the people on the risk management team to respond flexibly to situations covered in the plan. Nonetheless, the risk

management plan has to assign specific authority to the team to act on those situations that are the responsibility of the risk management process. The assignment of high-level roles and responsibilities also ensures that the routine supervisory and budgetary authority, which is needed to conduct the process as a normal part of doing business, is expressly assigned to the individual members of the team.

Finally, the concepts associated with risk management have to be defined in clear organizational-specific terms. That definition is necessary in order to align the organization's overall security objectives with its business objectives. In that respect, a comprehensive and detailed definition of key terms has to be provided as part of the planning setup process. The purpose of those definitions is to ensure a common vocabulary throughout the organization.

Definitions are important because most people's understanding of what constitutes risk is subjective. Consequently, it is recommended that the organization provides a precise specification of what constitutes a risk, the levels of acceptable risk, and the attendant approaches that will be used to address each risk. Specific directives for how to report risks and the thresholds for acting on risk reports also have to be preestablished for the various risk elements. The reporting requirements will also apply to active, residual, and accepted risks.

1.6.2 Identifying and Categorizing the Risk Environment

The next step in establishing effective risk management is to acquire comprehensive knowledge of the threat environment. That knowledge requires an all-inclusive record of the organization's assets, a statement of the acceptable levels of risk for each asset, and the constraints that will be placed on the protection of the asset by the available resources, technology, or existing policies. The outcome of the threat cataloging process is an alignment of the policies that will be used for risk management with the business goals of the organization.

That alignment is needed to conduct the trade-off process. Trade-offs will be used to decide the risk acceptance, risk avoidance, risk transference, or risk mitigation strategy that will be used to ensure each asset is addressed. When those trade-offs are planned, they should accurately reflect the organization's business objectives. An analysis of the priority of the information that enables the business objectives versus the threats to the information is necessary in order to decide where to invest the organization's security resources. Defining risk levels needs to be done with respect to their impact on the CIA of the data in the organization's operational systems.

Risk management coordinates three highly related factors within the operation, which are as follows: (1) the risks that can be associated with the organization's systems, (2) the business functions that are associated with the information in those systems, and (3) the extent of control necessary to manage each of those risks. The key to success lies in deploying the minimum number of controls to achieve

a desired level of assurance, given the intended purposes of each affected business function.

The risk control deployment process can be carried out in two different ways. The most common way to conduct the deployment is *ad hoc*. In the case of ad hoc risk control deployment, the controls are created to fulfill specific security needs. Those needs generally arise as a threat is identified. Many organizations use an ad hoc approach to risk management simply because the deployment of a coordinated set of controls is a difficult process to manage on a day-to-day basis. The ad hoc approach is cost-efficient because it only creates controls that are needed at the time. Nonetheless, it is almost certain to result in flawed protection because the organization is reacting to events that are occurring rather than deploying coordinated protection to prevent them from happening in the first place.

Another approach to risk management is the *coordinated approach*. Because it is meant to provide comprehensive protection, the coordinated approach offers more effective risk management. It deploys a series of risk mitigation baselines in a defense-in-depth scheme and is composed of a rational set of increasingly rigorous technical and behavioral controls. In most baselines, the electronic controls are automated while the behavioral controls entail a series of well-defined human-centered actions intended to produce a desired outcome. Each baseline is deployed to achieve specific risk management objectives and is prioritized in terms of the criticality of the data. Nevertheless, the creation, deployment, and ongoing monitoring of the baselines is both time-consuming and costly. Therefore, the degree of assurance justified under this scheme always has to be balanced against the level of effort and cost that is required to implement and maintain it. The aim of the coordinated approach is to deal only with the priority risks to the organization. In that respect, it takes active coordination to create and maintain an effective array of behaviors to manage the risks deemed most critical.

Because cost is a factor, a precise specification of the maximum degree of acceptable risk is a prerequisite to making a realistic plan. The specification of the maximum level of risk is necessary because much of real-world planning typically involves deciding what level of risk the organization is willing to accept. A decision about the degree of risk that the organization is willing to accept will lead to an assignment of priorities. Understanding the value of an item enables an explicit decision about its priority. The priorities then drive decisions about the practical form of the response. The value assigned is typically expressed as the level of acceptability of the risk. Consequently, acceptability is typically expressed in operational terms like, "Spend whatever it takes to ensure that this risk does not occur," all the way down to "The harm the risk would cause does not justify the cost of addressing it." Nonetheless, in order to decide about the level of risk, the decision-maker has to first know the value of the information the organization possesses.

Decisions about the acceptability of risk lead directly to a coordinated security response. Thus, the risk management process involves a technique that establishes a substantive, usually resource-based, link between every identified threat and the benefits of managing it. Operational factors that enter into that analysis include issues such as "What is the level of criticality of each particular information asset and what is the specific degree of resource involvement?" Therefore, threat/risk evaluations have to answer one key question at a minimum: "What is the trade-off between accepting the risk and the harm it can cause?"

1.6.3 Risk Assessment

The overall purpose of the risk management function is to maintain an appropriate set of risk controls. Therefore, ongoing assessments are a particularly critical part of that overall purpose. They are required because all control sets have to be periodically assessed in order to ensure that their protection is relevant and maintain their effectiveness. Risk assessments are important because they identify the specific threats to the organization, how likely those threats are to occur, and the consequences of each threat should it happen. Because knowing where risks lie is a fundamental precondition for managing them, the term "risk assessment" is sometimes used interchangeably with "risk management."

Moreover, risk assessment is not the same as risk management. Obviously, knowing the likelihood and impact of each potential threat is an essential precondition to managing it. Risk assessment is a tool that supports the larger risk management function, rather than an end in itself. Risk assessments underwrite the overall strategy that is used to deploy the risk management process. Risk assessments inform managers as to where to deploy the necessary reactive controls to respond to a risk. Risk assessments also monitor the effectiveness of those controls once they have been put in place. Thus, risk assessment maintains effective and up-to-date knowledge about the threat situation. And in many respects, risk assessment is an underlying prerequisite to the conduct of the risk management function. They are needed because a systematic risk assessment can specifically direct the maintenance of the controls that the organization has deployed to do substantive risk management. The targeted information ensures the most efficient use of security resources. Risk assessment is an information-gathering function that focuses on understanding the nature of all feasible threats. Risk assessment identifies and evaluates each relevant threat, determines the threat's potential impact, and itemizes the controls that will be needed to respond properly.

In that respect, risk assessments should always answer two distinct but highly related questions. The first is "What is the certainty of the risk?" The answer to that question is typically expressed as likelihood of occurrence. The second is "What is the anticipated impact?" The answer to that question is normally expressed as an estimate of the loss, harm, failure, or danger. Ideally, both of these questions can be answered in easily understood terms. Understandability and credibility are key

factors, because the results of the risk assessment will guide the deployment and subsequent conduct of the risk management process.

All risk assessments provide two specific pieces of knowledge: (1) the probability of occurrence and (2) the estimate of the consequences. There is a logical sequence to how these two questions should be approached. Practically speaking, the first consideration has to be likelihood, since a highly unlikely event might not be worth the cost of further consideration. However, it is the estimate of the consequences that truly shapes the form of the response. That is because there is never enough money to secure against every conceivable risk and so the potential harm that each risk represents always has to be balanced against the likelihood of its occurrence.

Therefore, the fundamental goal of the risk assessment process is to maximize the operational deployment of the organization's risk controls. Risk assessment accomplishes that purpose by identifying existing and potential threats with the greatest probability of occurrence and those which will cause the greatest degree of harm. The options these created are then arrayed in descending order of priority and addressed based on the resources that are available. Since all of the decisions about the tangible form of the risk management process will depend on getting the order of those priorities correct, it should be easy to see why a rigorous and accurate risk assessment process is so critical to the overall success of any risk management program.

Risk assessments are built around tangible evidence. The evidence is usually obtained by conducting interviews and documenting observations of both organizational and human behavior as well as auditing system logs and examining any other form of relevant technical or managerial records. Because the sources of data about risk are diverse, the collection process has to be systematic and coordinated. As a consequence, every risk assessment should embody a commonly accepted and repeatable methodology, which will produce concrete evidence that can be independently verified. The gathering, compilation, analysis, and verification of data about risk can be time-consuming and resource-intensive. So, in order to ensure the effectiveness and accuracy of any particular risk assessment, the practical scope of the inquiry has to be precisely defined and should be limited to a particular question, or problem.

Risk assessments typically target the various standard areas of threat—electronic, human, and physical. The insight gained from each assessment is then aggregated into a single comprehensive understanding of the total threat picture, which serves as the basis for deciding how each threat will be addressed. Operationally, it is perfectly acceptable to approach the understanding of risk in a highly focused and compartmentalized manner, as long as the organization understands that the results of any specific risk assessment characterize only a part of the problem. In fact, the need to paint a detailed and accurate picture of all conceivable threats almost always implies a series of specifically targeted, highly integrated risk assessments that take place over a defined period.

1.6.4 Designing for Effective Risk Management

1.6.4.1 Context

Every risk management process has to be designed to fit its particular environment. Environmental considerations are the factors that have to be understood in order to fit the risk management process into the overall operating circumstances of the organization. Accordingly, the design should describe all technical and environmental factors that might impact the risk management process.

In that respect, the design has to ensure that the process is correctly aligned with the environmental, sensitivity, and security requirements within the operational context of the organization. That is because the organizational context always dictates the risk management approach. For instance, there will be a different set of risk management procedures where the operational context is top secret or highly secure and requires very rigorous approaches, versus one where the context is more relaxed. As a result, the operational context in which the process functions has to be clearly understood in order to design a proper risk management approach.

1.6.4.2 Scope and Boundaries

Once the context about the scope or area of coverage is understood, the actual assurance has to be explicitly defined. The definition should be the result of a formal planning exercise. Formal planning is required because tangible organizational resources are involved. And failure to define an accurate and realistic scope for the risk management process could result in deficient protection and wasted resources. Therefore, distinctive and meaningful boundaries have to be established for the conduct of the risk management process. In particular, the logical interrelationships have to be understood between components, since the dependencies between the various elements that fall under the risk management process have to be factored into the assurance process. Or in practical terms, an activity that is linked to one outside of the scope of protection would represent a vulnerability. Since scope is always tied to the actual resources available, understanding which components will be a part of the risk management process and their actual interdependencies will allow the organization to be more realistic about what it will be able to protect.

1.6.4.3 Roles and Responsibilities

The definition of roles and responsibilities is a critical step in designing the risk management function since they tie both personnel and financial resources to the activities that will be performed. It is also important to explicitly clarify the duties that are associated with each of those roles. Otherwise, participants are likely to bring to the party their own assumptions about what they are supposed to do, which could result in important activities falling through the cracks.

Roles and responsibilities are created by designating accountability for performance of each security activity as well as all of the organizational reporting lines that are associated with each role. In that respect, if third parties or contractors are responsible for any aspect of risk management, the responsibilities and reporting lines of both the contractor and the organizational unit must be clearly defined.

1.6.4.4 Definition of Priorities

In addition to identifying and relating the various resource elements, each of these elements has to be categorized in terms of their general priority. Priority is directly related to the criticality of the resource. It is essential to be able to know the priority of each component in order to decide how many resources to commit to its protection.

The determination of priority is based on a simple understanding of the purpose of each element. The description of purpose should convey the general importance of the element in the overall operating environment. The description of purpose satisfies two operational goals. First, it allows managers to make informed assignments of priorities for the protected components. Second, it allows managers to coordinate the implementation and subsequent execution of the information assurance functions that are assigned to each component.

1.6.4.5 Sensitivity of the Information

It is essential to specify the sensitivity of each item of information within the system. That is because the sensitivity of the information determines the levels of CIA required. Thus, this specification provides the necessary basis for determining the extent and rigor of the controls. The specification also provides the basis for deploying the selected risk controls that will be used to secure each component. The specification should not just be guided by a consideration of technical standards and protocols. Minimally, the specification of the sensitivity should also consider the policies, laws, and any relevant constraints that might affect the CIA of information within the system.

The outcome of that specification should be a detailed recommendation of how the particular requirement will be addressed by a specific control. In addition, the recommendation for each control should provide a justification for why that particular approach was taken. The aim of that justification is to explain the type and relative importance of the protection needed. Each type of data and information processed by the system should be classified based on the severity of potential negative impacts on the organization and the degree to which the ability of the organization to perform its mission would be affected, should the information be compromised.

The sensitivity of information should be characterized based on the risks a compromise would represent. The highest risk would be associated with compromises that would adversely impact critical information, or which might result in loss of life, significant financial loss, threats to national security, or the inability of the organization to perform its primary mission.

Moderate risks would be those risks that might not compromise critical information but where the losses would still have business impacts. Low-risk items would be those risks where information might be lost but it would not be vital to organizational functioning.

1.6.5 Evaluating Candidates for Control

The threats that comprise the risk environment of an organization need to be understood before precise steps can be taken to manage them. Therefore, all known threats have to be identified, their relationships to each other understood, and the potential actions that they could take to cause harm have to be characterized. This can be evaluated and understood using the RMF stages. That understanding will let the organization describe in accurate terms the factors that threaten it and what those threats are likely to cause in terms of harm and their likelihood of occurrence. This understanding can then facilitate the development of precisely targeted controls for each threat.

Threat modeling is a structured method that is used to analyze risk-related data. A successful threat modeling process requires a lot of "creative" thinking, in that every conceivable threat should be put on the table and assessed. Threat modeling allows risk data to be modeled and subsequently communicated among team members. The major steps of threat modeling begin with a determination of the scope of protected space that the model corresponds to. Then threats that might impact the components of that space are enumerated and specific details as to the potential likelihood and impact of the threat are collected.

In order to ensure that the analysis is comprehensive, data flow diagrams or similar information flow diagrams such as unified modeling language (UML)-based use-case diagrams are employed to help visualize and describe the target space. These diagrams can be very helpful to ensure inclusive coverage. Descriptions of potential attack vectors and the impacts of each of the vectors on the protected space are used to think through and then describe the actual attack behavior. In that respect then, all potential attack vectors should be able to be described and examined from an adversary's point of view.

Subsequently, the implications of each threat must be analyzed. This analysis is typically based on assigning a criticality score. A standardized criticality score is an important part of the threat modeling process because it allows analysts to classify each identified threat in terms of its likelihood and potential harm. That classification can then lead to a priority ordering of known threats from most dangerous to least dangerous. The ordering will allow management to concentrate resources on

the threats that have the greatest potential for harm. It will also let managers assign fewer resources to lower priority threats. It is this classification process that allows managers to build logical and substantive defense-in-depth schemes.

A focus on priority differs from the typical low-hanging fruit approach. Nevertheless, the implementation process has to be based on some kind of quantitative or rational method for assigning priorities. Without priorities to guide the implementation, it is likely that the easiest to understand or most obvious threats will be addressed first. That approach would, in essence, disregard the business value of what was being protected. Given the requirement for thorough understanding in order to assign practical priorities, it is important to have a commonly agreed upon starting point to base the comparisons; this is the role of threat modeling. Threat modeling goes a long way toward putting quantitative and systematic implementation of the measures to control risk on a systematic and logical footing.

1.6.6 Implementing Risk Management Controls

The controls for risk management differ in their purpose and specificity. It is important to keep this difference in mind when designing and then assigning control activities because the people who will actually be executing each control need to know exactly how to perform all of the tasks that are necessary to make the control effective.

As a consequence, it is important to ensure that management types are not asked to perform highly technical tasks, just as it is equally critical that technical people are not asked to perform managerial activities. In both cases, there is the potential that the activities that underlie the control will be either misunderstood or misapplied. It is also important to understand the operational status of the control.

Knowing the existing operational status of the control, or even whether the control actually exists, is important in the design process. This is because some controls will already be present in the legacy scheme, while others will not have been created yet. Therefore, it is essential to have a complete understanding of where a procedure has already been implemented and where it has to be developed. This understanding is based on whether each necessary control item is operational and effective or not actually operational as originally planned.

It is common to have part of the control in place while other parts are still missing. If some parts of the control are implemented and others have only been planned, there should be an explicit specification of the parts of the control that are in place and the parts that are not. Where there are planned measures, this description should also include a list of resources required to make them operational and the expected timeline.

Finally, situations will exist where controls would be desirable, but it would be neither cost-effective nor feasible to implement them. If this is the case, then those controls should be noted for future planning as well as potential long-term monitoring of the risk that the measure was meant to manage.

1.6.6.1 Management Controls

Management controls are behavioral and based on policies designed to employ the organization's risk management procedures. Examples of management controls are incident response, security assessment, and planning controls. The nature of management controls is to manage risks through human-based actions rather than technology. These controls are typically designed based on a risk analysis, which should support a comparison between the costs of the applicable controls and the value of the information resource they are designed to protect.

Management controls are deployed based on the impact of the threats that they have been designed to address. It is important to design the appropriate administrative, physical, and personnel security controls into the risk management process from its inception. Because risks come in a number of forms, there can be an extensive range and variety of risk management controls.

Management controls are primarily enforced by the testing and review process. Therefore, the design must ensure that tests are performed during the development of the risk management process. The aim of those evaluations is to confirm that all of the necessary controls are an established part of the risk management process.

1.6.6.2 Technical Controls

Just as with the management process, the technical controls should also be well defined, understood, and followed. From a risk management standpoint, the most obvious technical controls are those that underlie the access control system. Technical controls are important and should be monitored closely. The monitoring of technical controls is an essential aspect of management accountability as well as a technical issue. As a consequence, the monitoring of technical controls from a managerial standpoint is often associated with audit procedures. A complete audit trail and a chronological record are evidence of adequate monitoring. The use of system log files to monitor system behavior is an example of this type of control.

1.6.6.3 Risk Type

Risks represent a threat to some aspect of organizational functioning. Moreover, the management of risk is a complex process with lots of inherent detail. As mentioned previously, in order to implement the risk management process, it is necessary to classify and understand the nature of the threats that are present in the organization's current operating environment. In general, threats can be classified into two categories, *known* and *unknown*.

Unknown threats, also known as *asymmetric* threats, are not predictable and not subject to management by standard risk management methods. Because of their unpredictability, they do not lend themselves to specific techniques for analysis. *Known* threats are those that should be logically expected to occur. Thus, another

name for known threat is *intrinsic risk*. In many cases, the probability of occurrence and subsequent impact of an intrinsic risk can be estimated. Intrinsic risks can be managed and minimized by an effective risk management program.

Accordingly, the organization has to adopt and follow some kind of structured process to identify, classify, and provide a meaningful response to the intrinsic risks that fall within the scope of the risk management process. This is the general aim of the RMF process. The RMF process can be employed to organize and coordinate the risk identification, analysis, and planning activities of a comprehensive risk management program.

Areas of intrinsic risk can be classified into three generic categories: management, operational, and technical. Using these categories in some form of checklist, managers can systematically work their way through a practical risk management situation and evaluate the status of each of the standard risk items on the list.

The management risk category encompasses the potential risks to the organization's information assets or documentation, as well as any of the risks that are associated with the assignment of roles and responsibilities and the risks represented by a failure to do proper contingency or configuration management planning. These are very large areas of organizational functioning and so their analysis requires extensive coordination. And because of the sheer scope of each of these areas, the analysis process itself usually requires a large number of participants. Managers can use the identification, assessment, and the select and implement stages of the RMF process as a roadmap to guide the deployment of the necessary controls to ensure a persistent risk response.

The second category includes the operational risks. These types of risk are much more focused and detailed. Operational risks involve threats to the operational environment that the organization has to manage, such as ensuring the identify management function and making certain that the identification and authentication processes, auditing, malicious code protection, long-term system maintenance, and communications security functions are properly ensured. These areas require the coordination of complex managerial and technical activities. Because of the complexity, the assurance of these areas has to be detailed and closely controlled. The RMF stages allow managers to both coordinate the threat identification effort and aggregate the huge amount of data that is normally collected in order to ensure that risk controls are effective and persistent.

Finally, there are the risks that are associated with the technical controls. Those include the predictable threats to electronic systems; however, they also include any electronic controls over media and the physical and personnel security environment. The technical risk category even includes risks that reside in the cybersecurity education, training, and awareness function. Because of their diversity and inherent complexity, every technical risk area has to be very well defined in order to be properly analyzed. A checklist of items for analysis is useful in facilitating this process and provides the necessary structure for the analysis. A checklist will also ensure that the right data is captured for each category and that the eventual analysis is appropriate.

1.6.7 Assessing the Effectiveness of Risk Controls

Forms of process assessment and measurement are important elements of good management practice. Assessment tells decision-makers whether or not their operational objectives are being met, that the results they are getting are in line with expectations, or even whether a process is under control. Risk management is no different than any other management activity in that regard. Good risk management requires appropriate measurement that accurately reflects the present threat picture of the organization. Nevertheless, proper assessment relies on the availability of meaningful standard measures.

Qualitative and quantitative measures can both be used for risk analysis. Both qualitative and quantitative measurements allow the organization to prioritize its risks and responses. The qualitative and quantitative measurement processes both assume that risks can be analyzed and that that analysis can be used to deploy the controls necessary to manage risk.

1.6.7.1 Qualitative Measurement

Qualitative measurement does not utilize actual metrics, but rather focuses on relative differences. Graphic scales are commonly used in qualitative analysis. Numbers may also be used, but they are merely markers for comparison value, not actual representative quantities. The end result of a qualitative risk assessment is a matrix of threats that differentiates between different relative levels of likelihood and impact.

In qualitative risk analysis, the measures that are used are descriptive, typically a set of nominal values such as high, medium, and low. These categories are then assigned numbers so that the weights of relationships can be characterized. Using those nominal values, it is possible to distinguish between items receiving a score of high versus those receiving a score of medium, for instance. However, it is not possible to truly rank different elements of the same class. So, the actual measurement itself is not precise. Nevertheless, since one of the main purposes of the risk analysis function is to determine priorities, qualitative analysis can be useful.

1.6.7.2 Quantitative Measurement

If there is a need for a more granular understanding of the risk situation, then quantitative analysis methods can be used. The value of quantitative methods depends upon the quality of the data being used. For instance, in the case of something like an actuarial estimate, hard evidence like the accuracy of records of birth and death and the causes of injury and loss, coupled with other factors, can be used to build predictive mathematical models. These models can be created and studied by analysts and the results from previous time periods can be compared with current results. In the case of risk management, accurate and reliable measures are difficult if not impossible to obtain while the changing nature

of the technology will restrict the application of time series studies. Therefore, in practice, a blend of both quantitative and qualitative measures is often used to arrive at the desired understanding.

1.6.8 Sustainment: Risk Assessment and Operational Evaluation of Change

Because the business environment is constantly changing, it is necessary to do continuous operational assessments of the risk environment in order to assure the validity of the risk management controls for the organization. Operational planning should be aligned with business goals and their accompanying strategies. The outcome of the assessment planning process must be a relevant monitoring of the current risk picture within business constraints.

All plans for any form of risk management process should be based on consistent standard assessment. Consistent assessment processes are important because management will use assessment data to make decisions about the degree of risk exposure as well as the types of controls that will have to be deployed. Accordingly, all of the metrics included in the risk evaluation process must be unambiguously defined in the plan. Those definitions can then be used to ensure that the data from the assessment process is consistent.

Consistency of measurement is a critical factor because stakeholders have to share a common understanding of the precise nature of the threats that the organization faces in order to trust the management response. As a result, it is important to make certain that there is reliable understanding of what a given assessment result means. If the various individuals who are involved in the risk management process interpret the information differently, there is a potential for uncoordinated and ineffective operational response. Additionally, there is the issue of credibility when it comes to the data itself. If there is no clear definition provided to function as the basis for measurement, then it is hard for decision-makers to rely on the data.

The activities that are involved in operational assessment are planned and implemented in the same way as other types of organizational assessment activities. That is, the operational risk assessment process employs risk evaluations to decide about the nature of emerging threats. Even so, rather than producing an overall risk management strategy, the goal of the operational risk assessment is to say with certainty that the currently deployed set of controls properly address the right threats. The assessment also seeks to prove that the controls continue to be effective given the overall aims of the business.

If the controls that are currently deployed do not address the aims of the business, then the operational risk assessment should provide all of the information necessary to allow decision-makers to make any changes that may be needed to achieve the desired state. Thus, any review report that contains recommendations

for change is typically passed along to the people who are responsible for maintaining the operational risk management process instead of the top-level planners who initially formulated the response. The aim of that report is to provide explicit advice about changes that must be made to the current risk management controls.

Planning for operational risk assessments involves the establishment of a standard schedule for each assessment as well as a defined process for problem reporting and corrective action. The routine nature of these reviews means that the organization should treat operational risk assessment exactly as it would any other continuous organizational process. That is, the process should be resourced and staffed to ensure that it functions as a part of the everyday business operation.

Operational risk assessment does not typically entail the sort of strategic planning focus that was involved in the formulation of the security strategy. Instead, it makes use of a defined set of performance criteria to evaluate the performance of the routine operation of the risk management function. Those criteria are typically laid down during the formulation of the initial risk management strategy. Consequently, every risk control that is deployed should have a clear set of standard criteria built into its specification.

These criteria should be both quantifiable and capable of being recorded and kept in a meaningful manner. Additionally, the assumptions about cost and occurrence that were part of the original decision to deploy each control should also be stated as a means of maintaining perspective on the operational intent of that control. The purpose of standard performance criteria is to allow decision-makers to judge whether a control is performing as desired and continues to achieve its intended purpose. The organization will use the data produced by the operational assessment process to ensure the effectiveness of its risk management scheme.

1.6.9 Evaluating the Overall Risk Management Function

The real proof of a risk management program's success lies in the operational outcomes of the controls that have been deployed for risk management. The test is whether the controls have achieved the desired business outcomes when it comes to risk mitigation. Control performance audits and assessments can be used to verify that the operational controls are functioning as designed and intended. Moreover, assessments can produce quantitative evidence that the control set is effectively controlling risk.

The assessment process itself is mainly a retrospective analysis of outcomes that is designed to verify through logs, record checks, and visual confirmation that the currently deployed control set has successfully covered the priority risks. The assessment examines the operation of those controls over some defined period in order to evaluate whether the organization is actually operating as planned. The assessment also attempts to characterize the effectiveness of each control based on the historical data that is recorded about its operation.

An audit adds a series of planned tests of the actual functioning of the process in order to confirm that its control features are functioning as they were designed to do. Both assessments and evaluations are designed to cover the entire breadth of the control set. Periodic audits are necessary for any organizational function. They are needed to ensure that the program is still meeting the objectives of the organization.

Risk management programs are no different in that respect. So, one of the important elements of the risk management process is the periodic execution of an audit that is designed to assess the overall effectiveness of the risk management program. Two types of audits are commonly used, a *time-based* audit and an *event-based* audit. It is generally a good idea to utilize both types of audits in practice, in order to ensure complete assurance.

A *time-based* audit is one that occurs at regular intervals, ranging typically from 1 to 3 years. These are top-down, comprehensive audits that are designed to examine all aspects of the risk management program against the business objectives that are currently in place. The purpose of time-based audits is to ensure that the risk management operation stays current with the business strategies and the ever-changing threat environment of the enterprise.

An *event-based* audit is much less comprehensive, but much more focused on a particular aspect of the risk management process. Like lessons-learned and after-action reviews, event-based audits are meant to capture and record information about a particular aspect of the risk management operation. For instance, if a business unit is reorganized, the business objectives may change. Because that change would represent a significant modification of the operating environment, it would be a good idea to make sure that the risk management program continues to support the goals of that unit. For the same reason, it is also important to audit the risk management situation after an actual incident has occurred in order to ensure that the outcomes of the incident reflect the desired results.

The objective of both of these kinds of audits is to ensure that the risk management program stays in step with changes in the business environment. Regardless of the type of audit that is conducted, there are some common elements that should be looked at as a part of each audit. The first of these elements are the controls themselves. In essence, the audit should determine how effective these controls were in detecting and responding to the threat that they were deployed to prevent. Additionally, the audit should confirm that there was not a need for additional controls for that particular incident.

In conjunction with the assessment of the actual control set, the audits should also examine the effectiveness of the policies and procedures that guide the implementation and routine operation of those controls. Those policies and procedures should be proven to align with the criteria for accepting the residual risk levels within the environment, as well as address the threat at the level of protection that is required. If the need to add additional controls, policies, and procedures, or

modify existing ones, is identified, then the audit report should itemize what those changes should be.

In addition to operational audits, a standard policy should be defined for conducting audits. As most organizations have an internal audit function, the audit of risk management processes and procedures should be built into their regular internal audit function. Conducting an audit of the risk management process as part of regular internal audit activity is an appropriate way to address the need for periodic audits of the risk management process. Rolling the assessment of the risk management function into regular internal audit activities is yet another way to institutionalize the risk management process.

1.7 Chapter Summary

In some respects, this book is as much about standardization as it is about risk management. Hence, this chapter presented an overview of the role of the standardization process in ensuring a consistent response to a given issue of importance. This includes a discussion of why information assets are difficult to protect as well as applying commonly acknowledged best practices to ensure an informed response. Specifically, we presented the issues involved in implementing a standard process including the benefits that derive from it as well as the potential pitfalls.

In practice, organizations design, implement, and follow some form of systematic process to establish a persistent operational risk management process. The design and management process is a strategic activity, in that it involves long-range considerations. Thus, planning for strategic risk management is necessary in order to ensure continuous risk assurance. And a formal strategic planning process is necessary to implement an organization-wide risk management process. Risk management itself must incorporate all of the elements of the business within its scope and the process should reach to the boundaries of the organization.

The outcome of the implementation of a risk management process is a concrete, organization-wide risk management scheme. The scheme will balance the aims of long-term risk control policy with real-world conditions and constraints. The atomic-level components of the risk management process are a set of substantive controls that ensure the requisite level of assurance against loss. These controls should be traceable directly to the policies that defined their need. This is a closed-loop process in that the ongoing alignment of risk controls to policies fine-tunes the evolution of the substantive risk management process and ensures its effectiveness in all operational settings.

One problem is that the term "risk management" is rather amorphous. So, the overall process itself requires a concrete statement of what risk management comprises. That statement is needed in order to make the practice standard. Standardization is important because a lack of effective, coordinated implementation and execution of the elements of the process has made overall risk management

efforts ineffective. One does not need to look any further than the increasing number of incidents in cyberspace to confirm that.

The other important justification for this standard is that the RMF also defines the basis for a comprehensive strategic governance approach to risk. A governance rather than technical approach is a highly advantageous strategy because, notwithstanding the issue of whether the cybersecurity function itself can ever fully embrace all of the issues associated with assurance, a governance-based solution is more easily understood and acceptable to the managers and nontechnical people who comprise the bulk of the organization.

This book starts from the assumption that a standardized risk management process should be applied corporation-wide. In that respect, risk management becomes a strategic issue rather than a narrow technical concern. The reason to adopt an organization-wide risk management approach is to avoid the dysfunctional effects of a typical piecemeal solution. The alternative approach to piecemeal is a formally defined and instantiated architecture of comprehensive risk management best practices, which are specifically aimed at optimizing risk controls across the company. As with any complex system, formal risk management practice can only be implemented through a rational and explicit planning process. The planning activity fits the strategic purposes and responsibilities of standards-based risk management to the security needs of the organization. From this standpoint, and throughout the rest of the book, it is the creation of that strategic risk management capability, which the RMF leverages, that will drive the presentation and discussion of the framework.

In simple terms, the risk management process assesses the likelihood that any given action will adversely impact something of value to any given entity. This includes such things of personal value as money, health, or even life. Once those risks are known, the risk management process deploys all of the measures that are necessary to ensure that consequent harm does not occur.

Because identification and understanding are important aspects of risk management, assessment provides the fundamental focus of the process. In essence, risk management is operationalized by a continuous process of assessing the organizational environment aimed at identifying and understanding all of the potential threats and the negative impacts that might affect the business. Once these have all been identified and characterized, then specific steps are devised and implemented to mitigate any adverse outcomes.

What is required to manage cybersecurity risks is a complete and provably effective framework that ensures the proper coordination and use of all appropriate methods in the execution of the process. The framework should be expected to consolidate provably correct approaches into a single logical and coherent model of operation. That model will contain all of the commonly accepted security best practices necessary to provide effective mitigation and management of all known risks to individuals, operations, and assets of the organization.

The NIST's RMF was designed to offer a structured yet flexible means for analyzing and deciding how to alleviate the risks that arise from the information systems within an organization. The idea of adopting a coordinated set of formal risk management practices is a relatively new concept. Cybersecurity risk encompasses all of the risks that are related to the use of ICT. Thus, the risk management approaches that are specified in the RMF are intentionally broad-based. That is because those recommendations are meant to dictate how to assess risk and employ the appropriate risk mitigation strategies for all conventional ICT organizations.

The requirement implies the need for a single umbrella model that defines the elements and relationships of the risk management process. The specific steps for risk management take place within the structure created by this overarching model. And these are captured in the appropriate supporting NIST and ISO security standards and guidelines that apply to that particular problem. The framework was derived from and builds on the collection of ISO, IEEE, and NIST standards. It also consolidates information from various standard body publications and provides examples of ways to implement those standards and guidelines.

An important feature of the RMF is the fact that it provides a practical basis for developing and maintaining comprehensive risk management controls for all aspects of a business's information assets. The objective of the RMF is to provide a common sense basis to develop, implement, and measure effective risk management practice. It is implemented through an organization-wide participative process and any business that has faced compliance issues with FISMA and NIST should be able to easily follow the RMF process.

Since the RMF touches on every aspect of how to assess and manage risk, it guides organizations through a step-by-step evaluation of their needs and responsibilities with respect to their ICT function. The process itself is generic and provides only the direction; it does not dictate the specific controls necessary to manage risk. Thus, the generic assessment and implementation approach must be adapted to fit every given situation.

In essence, an optimum approach is engineered out of the RMF model for each individual organization. The understanding of risk that the RMF provides and the appropriate set of control objectives selected from NIST SP 800-53 Revision 4 comprise the actual form of the eventual response. Accordingly, the approach to implementing the RMF is hierarchical. Or in essence, an explicit cybersecurity solution is evolved for any given unit, at any level of definition top-down within the reference model provided by the RMF.

And in that respect, the RMF assumes that specific cybersecurity approaches will be tailored to the outcomes of the common assessment process that is specified within the general framework of the RMF. This is accomplished in three steps. Once the threats, vulnerabilities, and weaknesses that the organization faces are assessed and their likelihood and impact are determined, policies are defined for each applicable control area. This serves as a foundation for tailoring.

The real proof of a risk management program's success lies in the operational outcomes of the controls that have been deployed for risk management. The test is whether the controls have achieved the desired business outcomes when it comes to risk mitigation. Control performance audits and assessments can be used to verify that the operational controls are functioning as designed and intended. Moreover, assessments can produce quantitative evidence that the control set is effectively controlling risk.

Glossary

best practice: a set of lessons learned, validated for successful execution of a given task

controls: a discrete set of human, or electronic behaviors, set to produce a given outcome

control performance: the operational results of control operation within a given environment

FISMA: the Federal Information Security Management Act of 2002

impact: a specific outcome or harm that might result as a consequence of a given threat

likelihood: the probability that a given event will occur, usually expressed as percent

NIST SP 800-53 Revision 4: the National Institute of Standards and Technology Security and Privacy Controls for Federal Information Systems

risk management: formal oversight and control of the threat mitigation actions of an organization

risk management scheme: specific architecture that embodies the overall strategy for risk mitigation

standard framework: a commonly accepted formal statement of best practice for a given topic

standardization: process of ensuring systematic common execution of a responsibility, or task

strategic governance: the overall long-term management control process of an organization, always administered from the top

strategic planning: the process of developing long-term plans of action aimed at furthering and enhancing organizational goals

systematic process: a process that has been standardized and embedded in the routine operation of the organization