CHAPTER

1

Information Security Threats and Risk

INTRODUCTION

Information Security Risk

This book is about estimating the vulnerability to unauthorized access to information by individuals with malicious intent. Attack scenarios range from simple visual observations of white boards or computer monitors and conversation overhears to sophisticated compromises of radiating electromagnetic signals. Many of these scenarios can be modeled using well-established physical principles that provide insights into the magnitude of the vulnerability to information security threats.

Such estimates may appear straightforward, but there are many scenarios of concern and the spectrum of vulnerabilities is broad. This is evident from the following examples:

- electromagnetic signals leaking from a computer located in a country known to sponsor information security attacks against foreign companies;
- a wireless network in the vicinity of a drive-by hacker with one of the network access points promiscuously radiating signal energy to the street;
- sensitive conversations that can be overheard by occupants of another floor in a multitenant office building;
- white boards, keyboards, and computer monitor screens in the direct line-of-sight of distant buildings;
- employees informally conversing on a company balcony while surrounded by properties of unknown control/ownership;
- information technology (IT) networks and systems that communicate via the Internet.

Traditional texts on information security and indeed most organizations often focus on the last bullet. In fact, entire departments are routinely dedicated to network security. It is no secret that the exploitation of IT vulnerabilities has increased in recent years, and the criticality of IT infrastructure demands a disproportionate share of attention. As a consequence, other attack vectors have been ignored despite the fact that they may be simpler to execute and could have equally significant impact.

The historical evidence suggests that IT risk deserve special attention. But despite that attention, traditional security risk management strategies have arguably been less than effective.

Information Security Science

Notably, these strategies include numerous and varied security controls. In fact, there is often an abundance of IT security data derived from these controls. But it is not clear such data are yielding insights into risk on a strategic level.

Indeed, IT risk managers are sometimes overwhelmed with data that are intended to identify the risk of information loss. But such data are traditionally used in support of tactical remediation efforts. Problems often recur because such remedies are inherently narrow in scope. Sometimes this abundance of data actually blinds organizations to the most significant risk factors for information compromise, which include business practices, security governance, physical security of information assets, and user behavior in addition to poor or inappropriate IT implementation.

One phenomenon that contributes to the ineffectiveness of current security strategies is the difficulty in quantifying IT risk. Why is this so difficult? The reasons are threefold:

- IT security incidents typically result from the confluence of related issues. The contributions of each issue can vary and it can be difficult to assess their relative magnitudes.
- Robust statistics on actual IT incidents are either nonexistent or not particularly helpful in assessing risk.
- Controlled experiments to determine the effectiveness of risk mitigation in IT environments are difficult to conduct.

The result is an absence of useful models pertaining to IT risk. So it is not easy to rigorously confirm the effectiveness of a particular security strategy.

The problem is exacerbated by the fact that IT protocols and systems have antagonistic objectives: ensuring data security and facilitating communication. To be sure, facilitating communication invites risk. In fact, the very existence of a network is a risk factor for information compromise. Despite continued efforts to ensure data security, these technologies spawn new vulnerabilities each day. The popularity of the Internet in conjunction with well-advertised attacks on systems drives the nearly singular focus on technology as both the culprit and the cure for information security ills.

Information security risk scenarios can admittedly be complex with interrelated elements. In many cases this complexity precludes the formulation of reliable risk models. Finally, the diversity of attackers and their respective motives makes precise statements on the likelihood of a future incident difficult if not impossible.

Information Security in a Routine Business Scenario

The following example illustrates the variety of information security issues that are relevant to even routine business scenarios. Consider an everyday meeting between individuals in a conference room. This event is likely repeated thousands of times each day around the world. Meeting attendees use the gifts provided by Mother Nature to generate, detect, and process acoustic energy in the form of speech.

What is the level of assurance that the acoustic energy will be confined to that conference room and not be overheard by individuals in other parts of the building? For example, acoustic energy propagating within the conference room might couple to building structural elements and be heard by individuals in an adjoining room or even another floor.

Moreover, if the conference room contains a device that enables individuals in remote locations to join the meeting, for example, devices manufactured by Polycom, the information security risk profile clearly changes. Telephones and telephone-enabled technologies are used to intentionally transmit acoustic energy beyond a particular room. Many of these devices are Internet Protocol (IP)–based, and are therefore potentially vulnerable to network-based attacks.

Any electronic device that transmits audible information must first convert acoustic energy into electromagnetic energy pursuant to long-distance transmission. The electromagnetic energy is transmitted via physical channels such as wire, optical fiber, and/or the atmosphere. However, a conversion to electromagnetic energy certainly does not confer immunity from signal compromise, but it does change the methods required to implement an attack.

Suppose further that a conference room targeted by an attacker contains a computer. The meeting organizer is using the computer to display confidential material to attendees in the room via a large monitor. He or she is also sending a PowerPoint presentation to remote attendees via the Internet. The computer is therefore being used to share the information with unseen audiences around the world. Visible displays of the presentation and electronic transmissions of data via the Internet offer adversaries a variety of attack vectors especially if an attacker owns the local communications infrastructure.

In addition, the conference room might be viewable to passers-by in the office hallway. Many modern conference rooms resemble large fishbowls that are located in well-trafficked office areas. These scenarios would provide opportunities for discreet information compromises. In addition, the same image that simultaneously appears on monitors around the world might be visible to passers-by in those distant venues as well.

White boards and computer monitors frequently display images within the direct line-of-sight of other buildings.[1] Such images can be viewed from significant distances with relatively inexpensive equipment. Even reflected images can be viewed in this manner given the appropriate physical conditions. Moreover, remote observation of visible information is much simpler than a network attack yet can yield information of comparable value.

This section concludes by asking the reader to now consider how the risk profile for the conference room scenario might change if someone with malicious intent had physical access to that room prior to a meeting. Many of the standard information security controls used to protect information would be undermined. The consequence is that now physical security controls could play an outsized role in ensuring information assets are secure from so-called "insider" threats. This scenario is an excellent example of the confluence of physical and information security risk, a recurring theme in this text.

Vulnerability to Signal Detection

Although quantitative estimates of information security risk are often desired, certain security scenarios are more amenable to such analyses than others. In particular, signal energy that is either intentionally or unintentionally radiated obeys well-known physical laws

[1]A visible image results from the reflection or direct transmission of energy within the visible portion of the electromagnetic spectrum. The intensity of background light, which lowers the contrast required to discern individual symbols, constitutes noise in this context.

that are formulated in terms of risk factors such as distance and time. Characterizing a threat in terms of these risk factors forms the basis for a quantitative model of vulnerability.

Let us be more specific about what is meant by a signal. A signal is a form of energy that is changed or "modulated" and thereby encoded with information. In other words, the process of modulation results in the transformation of mere energy into energy with information content.[2] A rigorous definition of energy is provided in Chapter 3. For now it is enough to know that both modulated and unmodulated energy obey the laws of physics. And it is modulated energy that is the target of attackers, and therefore provides the impetus to implement security controls.

Critical issues in estimating the vulnerability to signal compromise are the sources of ambient noise or interference and the physics of energy propagation within materials. Signals exist as either mechanical or electromagnetic energy, and each behaves quite differently depending on the material in which they propagate. The nature of these energy–matter interactions significantly affects the vulnerability to audible, visible, and electromagnetic signal compromises.

Environmental features in the path of propagation will affect the signal intensity and thereby impose detection limits as a function of distance. Developing a model for signal propagation that accounts for all risk-relevant features, that is, the risk factors, is critical to understanding the vulnerability component of risk for a given threat.

The physical nature of signals suggests that they should be detectable by some type of sensor. Furthermore, it would be reasonable to assume that these sensors are designed to detect specific forms of energy. In fact, it is the signal intensity that evokes a response by the sensor. Therefore the signal intensity in conjunction with the noise intensity is almost always the key to successfully estimating the vulnerability to information loss.

An attacker with physical proximity to a signal source could potentially detect and reconstruct that signal assuming the signal power is sufficiently greater than the ambient noise power within the signal bandwidth.[3] Moreover, signal energy is inherently promiscuous, and is agnostic to the identity of the individual attempting to detect it.

For example, if a Wi-Fi network is used to access the Internet, the signal is vulnerable to detection by both authorized users *and* an attacker sniffing for network traffic. A Wi-Fi signal carries risk if it is assumed to be undetectable in areas outside an organization's physical span of control and this assumption is part of a defensive strategy.

Although strong encryption (eg, WPA2-PSK) is now incorporated into the 802.11 wireless protocols to address this vulnerability, weaker forms of encryption are still prevalent (eg, WEP). In one ethical hacking exercise, my company exploited this exact vulnerability to gain access to a corporate network. In addition, unauthorized users can use other techniques to

[2]This statement is admittedly ambiguous due to differences in the meaning of "information." The colloquial use of the term refers to any data that convey meaning. The information theoretic interpretation, which will be investigated more thoroughly in Chapter 6, is the uncertainty or diversity associated with an alphabet or source of symbols. To illustrate the distinction, if I have a digital transmitter broadcasting all 1's, this is information in the colloquial sense. But the signal conveys no information in the information theoretic sense as explained in the same chapter. However, either interpretation works in this context if it is specified that modulation creates "content" rather than information.

[3]The effect of encryption is neglected here, which will not influence signal detection but is designed to thwart signal reconstruction if detected.

spoof the system and thereby connect to the network. Signal detection is a necessary precursor to signal reconstruction, so awareness of the risk factors that affect the vulnerability to detection should be top of mind.

Although quantitative estimates of vulnerability are useful, understanding the parameters that affect signal propagation at a high level is often sufficient. Qualitative insights can enable the identification of relevant security controls and thereby inform a risk management strategy. For those interested in more detail, references that offer more complete treatments of specific topics on security risk are provided throughout this text.

Assessing the Root Causes of Information Security Risk

A comprehensive assessment of the vulnerability to information compromise requires visibility into the root causes of information security risk. To that end, most texts on information security focus on technology when identifying issues requiring remediation.

An exclusive focus on technology ignores organizational issues that drive information security risk but are manifest as technology problems. In fact and as noted earlier, a virtual avalanche of information is generated by security technology controls. These controls can be effective in identifying tactical issues but the broader risk associated with security data so derived can be difficult to interpret pursuant to identifying systemic security problems.

Specifically and as mentioned previously, the principal sources of information security vulnerabilities in organizations are as follows:

- business practices
- a lack of security governance
- IT implementation
- the physical security of information assets
- user behavior

Business practices are inexorably linked to the organizational culture where the former is a by-product of the latter. In fact, it is not an exaggeration to say that the culture establishes the security posture in any organization, and reflects the outcome of the security–convenience dialectic that is being played out every day.

A lack of security governance is arguably driven by organizational culture as well. A lax approach to developing and enforcing information security policy follows from a culture of permissiveness or where creativity is encouraged on every level. In particular, a lack of well-designed security policies and standards correlates with a proliferation of information security vulnerabilities.

Although the previous discussion might seem to suggest otherwise, IT implementation is clearly a source of risk factors for information compromise. Poor implementation of technologies used to store, process, and/or transmit confidential information contributes to the overall risk profile of an organization, and must be addressed on both a tactical and a strategic level, but not to the exclusion of other sources of risk.

User behavior is deserving of scrutiny in identifying root causes of information security risk. User behavior in this context includes the history of websites visited, electronic access privileges (eg, Windows user, local administrator, domain administrator), physical access system privileges, history of internal resources accessed, and password complexity.

A user risk profile consisting of these risk factors enables a relative risk ranking of users, and thereby focuses monitoring and/or remediation efforts. Moreover, security perspectives at both the user and the device level provide a multidimensional view of risk, and thereby facilitate more effective risk management strategies.

The Physical Security of Information Assets

The physical security of information assets is relevant to the threat of information compromise and therefore represents an important source of vulnerabilities. Understanding the risk profile of data centers is particularly germane given the concentration of assets therein. The risk is driven principally by the use of virtualization and the related trend in using Cloud storage and computing resources.

Physical security is an area often neglected in traditional treatments of information security. This is unfortunate since the physical security strategy in data centers greatly affects the overall information security risk profile. The most obvious implication of a breakdown in a physical security device is that it increases the vulnerability of specific information assets to compromise. If that device malfunctions due to equipment failure or a network attack, those assets are less protected and therefore at an increased vulnerability to compromise.

In fact, physical security technologies can themselves be a vector for network attacks since most of these devices are IP-enabled and communicate via a local area network and/or the Internet. If a digital video recorder, network video recorder, card reader, or CCTV camera itself was compromised via an attack against its operating system or firmware, it could jeopardize that specific device as well as the network at large. The emphasis in this text will be on strategic security risk rather than the detailed workings of specific physical security technologies. Such technology has been addressed in many other texts [1].

Although nearly all organizations depend on networked devices to communicate and access confidential information, a surprising reliance on paper documents persists in many of these same organizations. The storage and transport of physical media containing confidential information presents surprisingly vexing security challenges, perhaps because organizations focus on network security and therefore neglect what appear to be more pedestrian threats.

Finally, network-based attack vectors against physical security devices and the details associated with remediation methods, for example, network isolation, application whitelisting, and signal encryption, will not be addressed in this book. This decision is again consistent with the fact that many other texts address such issues in exhaustive detail. Attempts at doing so here would potentially vitiate the treatment of security topics that are not typically discussed elsewhere yet can contribute significantly to an organization's information security risk profile.

The Likelihood Component of Information Security Risk

Measuring the vulnerability component of risk is necessary but not sufficient to develop a comprehensive view of information security risk. Assessing the likelihood of occurrence of a future threat incident clearly must be a factor in decisions on risk management. If a future incident is deemed unlikely relative to other threats, then resources might be better applied elsewhere.

The challenge is to evaluate the potential for incident occurrence if historical evidence of security incidents is rare or conditions vary significantly in time. This condition often reflects

reality. So how should a Chief Information Security Officer (CISO) or decision maker proceed in such circumstances?

There at least two methods to evaluate the likelihood component of information security risk: (1) perform statistical analyses of security incidents that relate to threat risk factors (this contrasts with attempting to count and analyze actual threat incidents, which, as noted earlier, is often not feasible) and (2) perform statistical analyses of threat incidents that can be modeled as random variables.

Risk factors will be discussed in detail later in this chapter, but the definition is introduced now given its relevance and importance: A risk factor for a specific threat is a feature that increases the magnitude of one or more components of risk for that threat. How are risk factors applicable to measuring the likelihood of a future information security threat incident?

In the absence of actual security incidents, analyzing incidents that relate to a threat risk factor offers a viable alternative. Since by definition a risk factor increases the likelihood, impact, or vulnerability to a threat incident, logic dictates that numerous incidents that relate to a risk factor are indicative of an increased potential and/or vulnerability to such an incident.

An example might be to analyze the number of incidents of unauthenticated access to restricted areas via piggybacking, etc. The successful circumvention of physical access controls to gain access to sensitive areas can yield relevant metrics on the quality of physical security even if such assets have not been compromised as a result. Successful password cracking is another example of where measuring a risk factor, for example, weak authentication, is indicative of the vulnerability to an actual incident.

If threat incidents are believed to occur randomly, it is possible to perform specific statistical analyses, and such a condition accounts for the second method of estimating the likelihood component of risk. If threat incidents can be legitimately considered random variables, well-understood statistical methods can be used to provide a quantitative estimate of the likelihood of occurrence.

It may seem ironic, but random processes confer a degree of certainty to inherently uncertain processes. This is because the standard deviation, which represents the uncertainty about the mean of a probability distribution, is specified for various distributions of random variables. For example, the probability that a given value selected from a normal distribution of values is within a standard deviation of the mean is proportional to the square root of the total population in the distribution.

Unfortunately or not, information security threat incidents that can be modeled as random variables are rare. Nevertheless, certain threat incidents might be amenable to such a model if only to provide crude estimates of risk. Chapter 13 details a method that enables estimates of vulnerability using this type of probabilistic approach. Therefore it can be helpful to be familiar with these methods and to apply them appropriately if judiciously.

INFORMATION SECURITY RISK

Understanding the distinction between a threat and a risk is a prerequisite for effectively communicating a risk management strategy. It is important because although threats and risk are closely related, they are not equivalent. Threats are the entities or conditions that cause harm, and therefore should be the focus of attention in a risk management strategy.

Evaluating the risk associated with a threat provides the impetus for going forward with security solutions as well as the requirements for those solutions. Security professionals should therefore address threats by evaluating the risk they present to their respective organizations. The following definition of a threat is fit-for-purpose, although there can arguably be many variations on a similar theme:

> A threat is any entity, action or condition that results in harm, loss, damage and/or a deterioration of existing conditions.

Given this definition, the spectrum of potential information security threats is quite broad. Threats to organizations might include thieves intent on stealing money, state-sponsored entities attempting to access company-proprietary or classified government information, and groups seeking to embarrass adversaries by exposing confidential information for political or economic gain.

It is this diversity of threats and their respective methods that drives the breadth of security risk mitigation measures. However, no organization can apply every possible mitigation method in equal measure without near-infinite resources. What is needed is a means of prioritizing threats in order to strategically apply remediation, which is precisely the point of a security risk assessment.

In that vein, a critically important role of the security professional is to identify the threats of highest concern (read: highest "risk"). This activity should be followed by measures that reduce his or her organization's vulnerability to those threats within the constraints imposed by budgets. Indeed, it is the finiteness of available resources that makes prioritization of remediation efforts a necessity.

So now that threats have been defined more precisely, what exactly is risk? All threats are described by a fundamental characteristic called risk, which is a set of three components as follows:

- *the impact or importance of a threat incident*
- *the likelihood or potential of a future threat incident*
- *the vulnerability or potential loss due to a threat incident*

These components collectively define the risk associated with a threat. In fact, risk can be notionally represented by an "equation" that is expressed as a product of the individual components as follows:

$$\text{Risk (threat)} = \text{impact} \times \text{likelihood} \times \text{vulnerability} \tag{1.1}$$

(1.1) should be read as, "The risk associated with a given threat equals the product of its impact, the likelihood of its occurrence, and the vulnerability to loss or damage."

For now, suffice it to say that assessing the magnitude of the vulnerability component of risk, that is, the loss, damage, or exposure to a threat incident, is the basis for many of the analyses in this book.

Importantly, the risk associated with a threat is not immutable, and the magnitude of each component can vary significantly depending on circumstances. Context is crucial in assessing risk. In fact, a security assessment is merely an abstraction without context. If one were to provide a high-level if formal job description of a security professional, it is to evaluate the

risk associated with the spectrum of distinct and impactful threats in light of scenario-specific parameters.

Identifying the spectrum of distinct and impactful threats is the progenitor of every security strategy. This task sounds simple, but determining what constitutes an impactful threat can be quite subjective and even controversial.

For example, some might argue that religion and television represent dangers to society. Yet many individuals, even intelligent ones, believe quite the opposite. With respect to distinctness threats that are seemingly different can actually be functionally equivalent in terms of the required risk mitigation. However, there is a test for distinctness that will be explained in the discussion on risk factors.

Analogies with the medical profession are often useful when thinking about concepts in security. Security threats are equivalent to diseases in medicine, and risk mitigation measures are analogous to therapies. Most reasonable people would agree that diseases make people worse off. So unless you are a bit sadistic, hearing that a relative, friend, or associate is afflicted with a disease would probably be unwelcome news.

In medicine identifying the need for risk management is usually relatively easy. Patients display symptoms that are manifestations of some condition. Remedies are sometimes prescribed as a prophylactic measure based on one's exposure to a microorganism, a genetic predisposition to an ailment, or some *risk factor* for a particular disease.

Once a disease or precondition has been identified, patients pay physicians (and insurance companies) to prescribe therapies. Such therapies often take the form of a drug. The effectiveness of that therapy will of course depend on the correctness of the diagnosis, but will also relate to each individual's physiological makeup since no two people are identical.

But fortunately people are biologically similar, or at least similar enough, and that fact is the key to the large-scale effectiveness of many therapies. If one believes otherwise, there should be a separate anatomy and physiology textbook for each person on earth.

Experiments can be conducted that leverage the similarity of humans such that the action of a specific therapy can be isolated from other variables, and thereby lead to a conclusion on cause and effect. The process leading to the approval of a new drug, which includes testing hypotheses on effectiveness, is typically quite protracted, and expensive.

First, experiments are conducted on animal models that use a control group to isolate the effect of a single variable, namely the drug in question. Researchers attempt to establish a causal link between the disease and the palliative effects of the drug while observing potential side effects. The type of animal is chosen because their physiological response can be extrapolated to humans.

Once the animal studies have concluded, and it is clear that the drug had the intended result without obvious harmful side effects, human trials can commence. So-called "double-blind" experiments are designed to eliminate bias where a statistically significant trial population is divided into control and test groups.[4]

Following the human trials and assuming a positive outcome, the drug is approved for general use by the Federal Drug Administration (US). As an aside, the average cost of research and development for a prescription drug is estimated to be $2.558 billion [2]. The point

[4]Double blind means that the identities of both the control and experimental groups are unknown to the study participants.

is that medical threat scenarios benefit from significant testing of hyptotheses relating cause and effect.

Contrast this with security scenarios. In general, threat incidents are relatively rare, and, importantly, there is often considerable variation in conditions that undermines the ability to isolate a variable under test.

One can simulate attacks on networks and applications. That is the point of conducting penetration tests. Such simulations will provide a degree of confidence in the resilience of specific security controls. But this is not the IT equivalent of a drug that confers broad immunity. The operational model, which consists of the user environment, is too complex, ephemeral, and varied.

INFORMATION SECURITY RISK ASSESSMENTS

In general, comprehensive assessments of information security risk are required to establish a thorough understanding of the risk factors affecting an organization. Furthermore, such assessments must be made with respect to risk-based policies and standards in the absence of useful statistics on incidents. Adopting a process to rigorously assess the risk associated with information security threats is essential to developing a coherent information security risk management strategy [1,3].

Since the essence of security is to mitigate the effect of threats, all estimates of risk should begin with identifying the spectrum of distinct threats. Threats were defined previously, but what is meant by "distinct" in this context?

Distinctness implies a set of characteristics that distinguishes one threat from another. Characterizing threats under general headings such as "terrorism," "street crime," and "hate crime" may be useful for sociologists and politicians, but it is not particularly helpful in developing a risk management strategy. So how *does* one specify that a given threat is distinct from another and why does it matter to a risk assessment strategy? These questions will be answered following a brief digression on risk.

Recall (1.1) was introduced as an operational definition of risk and was formulated in terms of three components, likelihood, vulnerability and impact. This was somewhat hyperbolically referred to as the Fundamental Expression of Risk. However, it is not a true mathematical equation because each component in (1.1) appears to have equal magnitude and this condition is not true in general.

One important feature to notice about this expression is that if a single component is zero, there is no risk. The implication is that if there is no risk, the threat being evaluated does not exist for all practical purposes. Put another way, absent one or more components of risk, a given threat is simply not *threatening*.

In addition, the notion of "cost" broadly defined is missing from (1.1). Although cost is not a fundamental component of risk *per se*, it plays an important role in real-world decisions on security.

For example, it is not uncommon to encounter security risk scenarios where the magnitude of one component of risk is significant but remediation is cost prohibitive. Therefore, despite the assessed risk no action is taken to address it. The cost associated with risk mitigation is a reality associated with real-world risk management processes that would not appear in a strictly academic view.

Although a measurement of risk is ideal, it is not always possible to provide a quantitative estimate. The reality is that a qualitative view of each component is sometimes the best option available. The good news is that such a view is often sufficient to make a meaningful decision on risk mitigation. Moreover, a sophisticated security risk manager understands when quantitative measurements of risk will yield meaningful results and when it is futile to even try.

With that background, the risk assessment process can now be described, and, in particular, the critical role of risk factors in developing an effective risk management strategy. As noted earlier, the first step in a security risk assessment is to identify the spectrum of impactful and distinct threats to an organization. In order to address the question of threat distinctness, the crucially important concept of a "risk factor" must be reintroduced and defined as follows:

> *A risk factor is a feature, characteristic or condition that enhances one or more components of risk for a specific threat or mode of threat implementation. It is the spectrum of risk factors that drive the required mitigation methods.*

The logic associated with risk factors as the basis for risk management is compelling to the point of appearing circular: If risk factors are those features that enhance one or more components of risk for a given threat, then addressing all the risk factors is required in order to effectively manage that threat.

A medical analogy is again illustrative. Consider the threat of cardiovascular disease. Some well-known risk factors for this threat are high blood pressure, obesity, a high concentration of certain types of cholesterol in the blood, smoking, lack of exercise, being male (or a postmenopausal female), diabetes, and a family history of cardiovascular disease.

These risk factors were determined through large population studies that enabled scientists to correlate the presence of a risk factor with the likelihood of a future threat incident. In other words, people had varying rates of heart attacks based on the number and magnitude of one or more risk factors.

The likelihood of a future threat incident increases by some quantifiable amount with each additional risk factor, an artifact of the plethora of data established over years of studying relatively homogeneous models such as humans. In other words, the more risk factors displayed by a patient, the higher is the likelihood he or she will suffer a heart attack in a specific interval of time.

The risk increases with the duration of the time interval under consideration.[5] An individual who displays all of the significant risk factors would likely be a candidate for aggressive medical therapy as determined by a bona fide medical risk manager, for example, a cardiologist.

A Venn diagram can be used to illustrate the intersection of risk factors, a condition that would amplify the likelihood component of risk for the threat of heart attacks as shown in Fig. 1.1.

[5]One way to think about the increasing risk associated with lengthening time intervals is to consider the limiting cases of time intervals, that is, $t = 0$ and $t =$ infinity. The likelihood of an event in an interval of 0 s must be zero. At the other extreme, in an infinite time interval the probability of an event would approach unity. For intermediate intervals the probability of an event is proportionate to the duration of that interval. However, recognize that this relationship between likelihood of occurrence and time interval duration might not be linear since the cumulative effect of risk factors could be exponentially increasing with time.

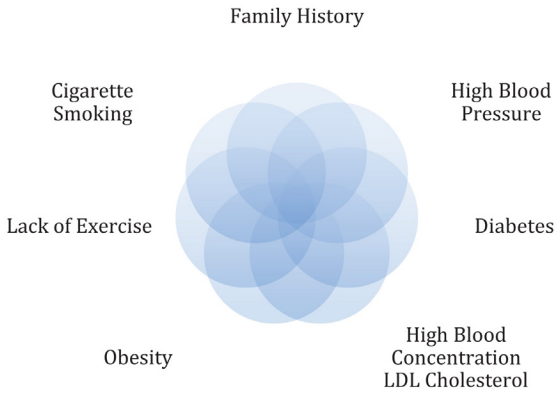

FIGURE 1.1 Intersection of risk factors for the threat of cardiovascular disease.

A similar diagram can be created for any threat. Physical security threats are illustrative of the utility of such diagrams. Consider the threat of vehicle-borne explosive attacks by anti-Western elements against the headquarters of an international bank. Risk factors for this attack might include the following:

- the country where the facility is located;
- the iconic status of this particular facility or the bank in general (in other words, a symbolic association with Western culture and/or a particular government);
- the historical use of this mode of attack by groups of concern;
- the proximity of the facility to vehicular traffic.

Note that the first three risk factors enhance the likelihood component of risk for this threat while the last one enhances the vulnerability component of risk. Understanding the nature of the contribution to risk for a given risk factor is important in managing the risk associated with each impactful and distinct threat. For example, reducing the profile of a company or facility would affect the potential for attack, but would do nothing to reduce the vulnerability or the potential damage/loss should an attack occur.

Fig. 1.2 illustrates the Venn diagram for the set of risk factors associated with a given target and relative to this threat. If all of these risk factors existed for a given target, the risk is enhanced relative to a target that possessed less risk factors.

To further illustrate this important point, if the impactful threats were groups concerned about the global hegemony of fast food corporations, the likelihood component of risk might be significantly altered from the anti-Western terrorists noted earlier. In that case the security strategy might not include this threat as a priority for remediation.

The long-awaited answer to the question of what makes one threat distinct from another can now be presented. Simply put, any two threats are equivalent if the type and magnitude of their respective risk factors are identical. Conversely, if their risk factors differ in either type or magnitude, the two threats are distinct and each threat must be addressed separately as part of a risk mitigation strategy.

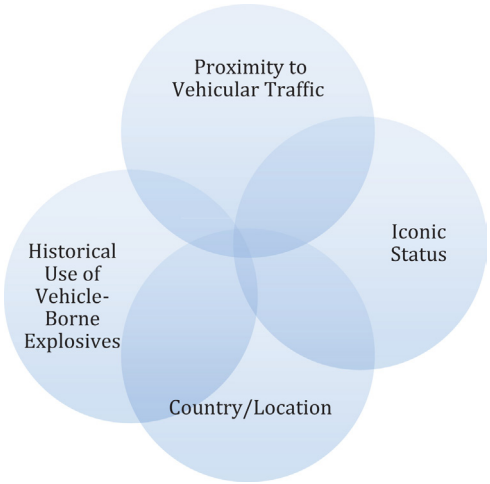

FIGURE 1.2 Risk factors for vehicle-borne explosive attacks by anti-Western groups.

This test for distinctness has a very practical implication. Namely, threats can be logically grouped according to their risk factors. In addition, simultaneously addressing the risk factors will effectively manage all of the threats with risk factors in common. Note that if one risk factor is not addressed, it means at least one vulnerability exists for each threat to which that risk factor applies.

The key to an effective risk mitigation strategy is to address all the risk factors for each distinct and impactful threat. A graphic that depicts the risk management process is captured in Fig. 1.3 [3].

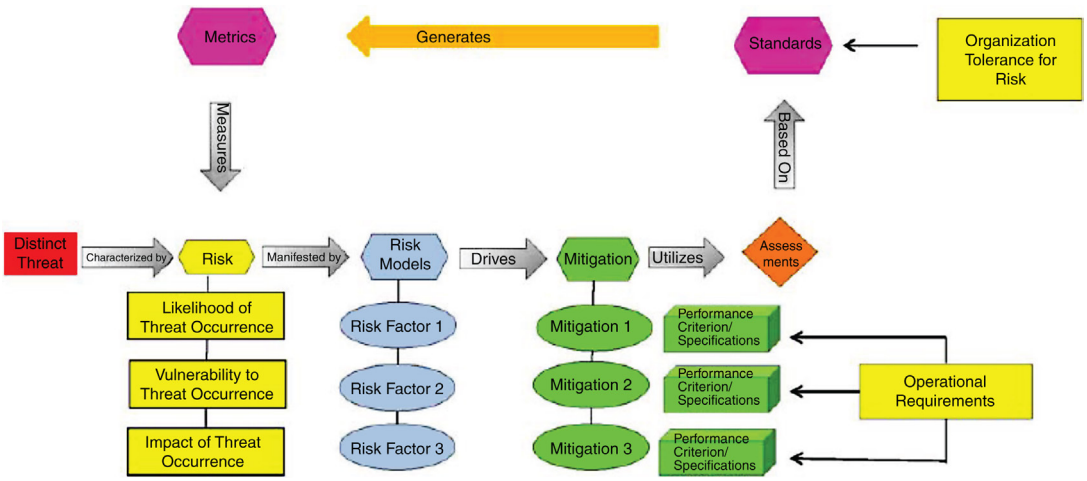

FIGURE 1.3 The security risk management process.

ORGANIZING INFORMATION SECURITY RISK ASSESSMENTS

Organizing an information security risk problem is often a useful initial step in assessing risk. In particular, establishing and then analyzing categories with common features facilitates coherent analyses. Furthermore, defining individual units according to specific features that are categorized in a hierarchy with descending levels of granularity can reveal patterns or themes. These patterns enable general conclusions about the organization as a whole.

Many descriptive sciences organize information this way. For example, biologists have created a hierarchy for all living organisms as follows: kingdom, phylum, class, order, family, genus, and species.[6] This method has enabled scientists to identify evolutionary trends. Presumably anyone reading this book belongs to *Homo sapien* the human genus and species, respectively.

Since organizing a problem according to specific features has applicability to many disciplines, formal methods have been designed for this purpose. Specifically, a so-called *dichotomous key* delineates a group of "things" according to a hierarchy of functionality or some other feature/characteristic. Why is this approach helpful to security risk analyses?

Organizing threats according to a hierarchy of features they have in common facilitates the identification of common forms of mitigation. Such an approach enables the development of a comprehensive and coherent risk management strategy since one can associate general modes of attack with specific mitigation measures.

Following this alignment of general features with mitigation, additional details might drive specific mitigation requirements that deviate from the general case. Unfortunately, sometimes the tendency is to initially focus on details, which is more likely to result in missing the big picture, which in this case translates to missing a particular risk factor.

There are a number of ways to construct a dichotomous key for security threats. One might begin by identifying the spectrum of attack vectors for a given threat. For example, if terrorism is the general threat of concern, one might first want to specify a hierarchy of attacks. Again, the purpose of such a hierarchy is to highlight common risk factors that are addressed by the same mitigation measure. A graphic for one version of a dichotomous key for information security attacks is shown in Fig. 1.4.

Fig. 1.4 is not an exhaustive exposition of the possible modes of attack. It is intended to illustrate one of many possible organizational schemas. Other versions could be constructed according to different organizational criteria.

Although not a dichotomous key *per se*, the NIST Cybersecurity Framework is illustrative of the benefits of a hierarchical structure in addressing complex security risk problems [4].

The NIST Cybersecurity Framework defines Functions, Categories, and Subcategories in a hierarchy of security controls.

It also defines four Tiers that correspond to increasing levels of sophistication: Partial, Risk-Informed, Repeatable, and Adaptive.

The five Functions or high-level security controls that are evaluated to establish a Tier rating are as follows: Identify, Protect, Detect, Respond, and Recover. These are admittedly too coarse to be actionable. However, this 50,000-ft. view is useful as a means of organizing a

[6]The phrase "King Philip came over for good soup" is a pneumonic used to assist in recalling the hierarchy of organisms.

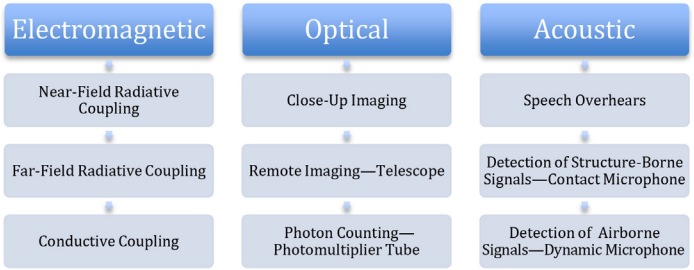

FIGURE 1.4 **Abbreviated dichotomous key for information security attacks.**

risk-based information security strategy. Metrics to rate specific NIST Tiers are suggested in Chapter 12.

Importantly, the NIST Cybersecurity Framework is not prescriptive, and therefore does not dictate requirements for specific controls. Equally if not more importantly it is not a checklist, which is what makes it adaptable to risk-based security assessments. Rather, it is intended to facilitate assessments of the processes required to assess information security risk *in context*.

Therefore, it enables risk-based decisions based on an organization's sophistication in implementing security controls rather than merely checking that such controls exist. It also facilitates addressing threats in a consistent manner across business units.

For example, one well-known information security threat is the covert exfiltration of confidential information via an IT network. Using the NIST Framework it is a relatively simple process to evaluate the maturity of a security strategy with respect to this threat.

Table 1.1 shows an indicative if high-level and incomplete analysis of this threat using the NIST Framework. Analyses of subcategories/controls relative to NIST Tier criteria are missing. But it illustrates how a NIST-based risk assessment might be organized.

TABLE 1.1 Abbreviated Assessment of the Vulnerability to Covert Data Exfiltration Using the NIST Cybersecurity Framework

Function	Risk-Relevant Category	Analysis of Risk-Relevant Subcategories/ Controls	NIST Tier
Identify	Risk Assessments	Virtual Network Model/Simulation and External Penetration Test Results	Partial
Protect	Network Segregation	Firewall Rule Set	Risk-Informed
Detect	Automated Tools	Data Leakage Prevention (DLP)	Repeatable
Respond	Incident Response (IR)	IR Team, Plan, and Exercises	Adaptive
Recover	Recovery Planning, Communications	Assignment of Recovery Team Roles and Responsibilities, Established and Redundant Communication Protocols/ Channels	Risk-Informed
Overall Assessment/ Tier			Risk-Informed

This exercise may not seem particularly fruitful for a single threat and a small organization. But security professionals are often required to assess threats and the risk they present to organizations that have global footprints and/or a large number of business units.

Furthermore, developing this type of framework can be invaluable in identifying risk-relevant "themes" that emerge from disparate data sources and across stove-piped business units. Once such a framework exists, it is a relatively short journey to identify a coherent, enterprise-level risk mitigation strategy.

There are at least four critical risk factors that must be explicitly delineated within a security risk framework as follows: (1) the sensitive/confidential information (affects the impact component of risk), (2) the business units that use that information (affects the vulnerability component of risk), (3) the modes of information usage/management/access by each unit (affects the likelihood and vulnerability component of risk), and (4) where the sensitive/confidential information exists within the organization (affects the vulnerability component of risk).

Table 1.2 captures these risk factors for a generic academic organization to illustrate the basic assessment methodology.

TABLE 1.2 Analysis of Risk Factors for the Threat of Information Compromise

	Human Resources	Finance Department	Records Management	Fund Raising	Health Department
Sources of Confidential Information	Employee Records (PII)	Social Security Account Numbers	Student Records (PII)	Donor Information	Patient Information
Key Information Storage Resources	File Shares, Central Application	File Shares, Central Application	File Shares, MySQL Database	Bespoke Application	File Shares
Mode of Information Storage and Transmission	Email	Email	Email	Email	Email
Mode of Information Destruction	Manual Purging	Manual Purging	N/A	Manual Purging	Manual Purging
Document Management	Office Storage	Office Storage	Central Archive Storage	Office Storage, Manual Destruction	Office Storage, Manual Destruction
Physical Security Risk Factors	Physical Keys with No Access History	Physical Keys with No Access History	N/A	Physical Keys with No Access History	Physical Keys with No Access History
Technology Risk Factors/ Vulnerabilities	Low-Entropy Passwords for Key Information Assets; Open Network Access to Data Sources	Low-Entropy Passwords, IT Equipment Not Centrally Managed; Open Network Access to Data Sources	Low-Entropy Passwords, Noncurrent Version of Databases; Open Network Access to Data Sources	Low-Entropy Passwords, Extranet Connections; Open Network Access to Data Sources	Low-Entropy Passwords, Extranet Connections; Open Network Access to Data Sources
Specific Authentication, Authorization, and Access Privilege Risk Factors	Manual Assignment and Removal of Access Privileges for File Shares	Manual Assignment and Removal of Access Privileges for File Shares	Manual Assignment and Removal of Access Privileges for File Shares	Manual Assignment and Removal of Access Privileges for File Shares	Manual Assignment and Removal of Access Privileges for File Shares

High-level security themes that emerge from Table 1.2 are as follows:

1. There is a concentration of information security risk. In other words, a small number of applications are used to manage confidential information across the enterprise.
2. An inherently insecure method (email) is used to transmit confidential information across the enterprise and outside the organization.
3. Manual processes are used to facilitate access privilege assignment and removal for file shares.
4. Low-entropy passwords exist. In other words, a limited diversity of possible password constructions for critical assets exists.
5. Mechanical locks and keys are used to secure paper documents. No physical access history is available as a result.

Once these themes on risk are specified, it becomes easier to determine a comprehensive strategy for remediation. In addition, identifying other risk factors provides justification for implementing compensating controls. In the example discussed earlier one might consider using two factors as a means of authentication to access the most critical applications if other vulnerabilities cannot be addressed.

Importantly, intersections of risk factors are apparent from this analysis thereby establishing priorities for remediation. For example, file shares that (1) contain high-impact information, (2) require weak passwords for authentication, and (3) are accessible from the Internet would likely be a priority for remediation.

Any attacker who gained internal access to the network, often accomplished via social engineering, would have little difficulty accessing high-impact information assets and subsequently exfiltrating the information contained therein.

This type of analysis can be done for any security organization. The key is to be precise about the threats, identify the critical information assets and their locations, and determine the spectrum of risk factors that enhance the vulnerability to the relevant attack vectors.

GENERAL RISK FACTORS FOR THE COMPROMISE OF SIGNALS

Information security threats can exploit vulnerabilities in IT protocols, intercept signals with encoded information, or steal a physical object that stores information. This section will focus on general threats and risk factors associated with threats to signals.

Estimating the magnitude of vulnerability for information security threats can be complex, and will vary according to the specific threat under evaluation. However, simple if approximate models exist for electromagnetic and acoustic signals because they obey well-understood physical principles. Therefore, quantitative estimates of vulnerability can be made and thereby enable mitigation strategies.

Moreover, developing a model of signal behavior and the associated risk facilitates analyses of a spectrum of threat scenarios. The result enables more fulsome estimates of vulnerability as well as the general effectiveness of risk mitigation. For example, establishing a model for signal intensity in terms of the distance from a radiating source enables an estimate of the vulnerability to signal interception by an attacker at *any* location.

Five risk factors affect the vulnerability component of risk for the compromise of signals:

1. the form of signal energy (electromagnetic or mechanical);
2. the intervening materials between the radiating source and the point of detection;
3. the physical proximity of a radiating source and the point of detection;
4. the signal bandwidth and the magnitude of ambient noise across that bandwidth;
5. the sophistication of the attacker.

Consider risk factor number three. If the vulnerability to unauthorized signal detection happens to vary strongly with distance from a radiating source of signal energy, it would be prudent to maximize the distance between potential adversaries and the signal source in accordance with this model.

Such insights are admittedly not particularly profound. The real challenge is to determine at what distance the signal becomes invulnerable to detection. Understanding the precise dependence of signal intensity on distance plus the magnitude of the ambient noise power enables such a determination. This and similar estimates represent the essence of the information provided in this book.

A simple graphic depicting the high-level risk factors for threats to radiated signals is shown in Fig. 1.5. More detailed risk factors for the compromise of radiating signals are presented in Chapter 6.

Finally, it is useful to create a holistic view of an information security risk assessment that tells the complete risk story from threats to remediation. It is also helpful to visualize the process depicted in this linear view since the issues are somewhat self-explanatory when presented this way, and therefore helps facilitate decisions on mitigation.

Table 1.3 illustrates this format for an abbreviated assessment. Note that it is organized according to (1.1). In addition, the cost of remediation is included, which facilitates security decisions.

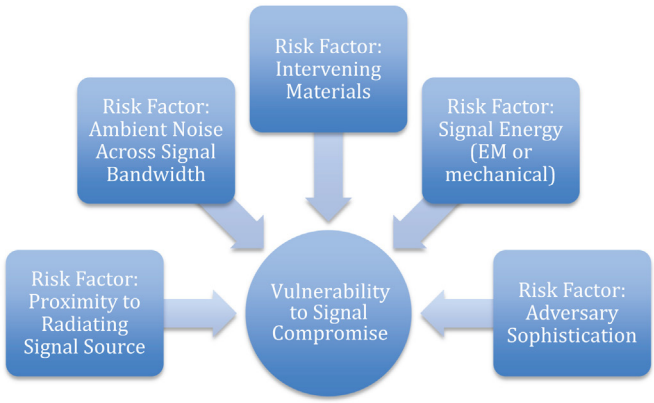

FIGURE 1.5 **Risk factors for threats to radiated signals.**

TABLE 1.3 Risk Assessment Summary Table

Principal Information Security Threat to the Organization	Risk Factors for Likelihood	Risk Factors for Vulnerability	Threat Impact	Remediation and Cost
Hacktivists seeking to steal personally identifying information (PII) for identity theft	Significant Internet Presence and Links to High-Profile Individuals	1. Open Network Architecture 2. Weak Authentication for Critical Systems	Information Compromise Leading to Reputational Damage and/or Regulatory Fines	1. Internal Firewall and Routing Tables 2. Two-Factor Authentication on Critical Systems Estimated Cost = $1 Million

Risk (threat) = likelihood × vulnerability × impact cost of remediation.

ESTIMATING THE LIKELIHOOD COMPONENT OF RISK

This book admittedly focuses on the vulnerability component of security risk to the exclusion of the other two components, likelihood and impact. This somewhat parochial view is not intended to trivialize the importance of these other components. In fact, assessing the impact associated with the compromise of information should be the first step in determining a set of proportionate security controls. Nonimpactful threats should generate minimal follow-up. Highly unlikely threats might warrant similar levels of inattention.

To be precise, the vulnerabilities analyzed in this text do not cause an incident. For example, it would be technically incorrect to claim that poor physical security controls at a facility increase the likelihood of a physical attack.

However, it would be correct to say that if an attack does occur, such vulnerabilities enhance the likelihood of its success and resulting loss/damage. Recognize that the likelihood of a future incident, which is the type of likelihood specified by (1.1), and the likelihood that such an incident will be successful are two different phenomena.

The impact component of risk is by definition organization-specific. Intellectual property and other proprietary information, personnel records, company strategies, etc., will vary based on each organization's mission, objectives, composition, and/or legal/regulatory requirements. Therefore, it is not particularly useful to discuss this component at length except to point out the overall criticality of determining if such information exists, where it is located, and the impact to the organization were it to be compromised.

On the other hand, the likelihood component of risk warrants additional discussion. The concept of likelihood is not well understood by security professionals and is a perennial source of confusion, which further motivates the following discussion. One theory why this misunderstanding exists is that it stems from the colloquial use of the term "likelihood" coupled with a misunderstanding of probability and statistics.

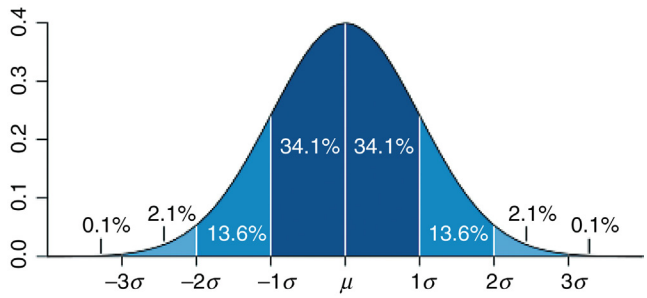

FIGURE 1.6 **The normal or Gaussian distribution.**

There is a very specific condition that must be satisfied in order to make precise statements about the likelihood of a future security incident. Namely, the incident in question must be a random variable and the outcomes are distributed according to a stochastic (probabilistic) process.

Ironically, the implication of being a random variable is that although the probability of guessing the value of a future event is completely unpredictable, the probability of guessing the value of one event from a distribution of events is quantifiable. Moreover, the predictability improves for larger numbers of events. This phenomenon will be explained later.

Qualitatively, a random event is where the value of the outcome of that event cannot be determined in advance. Yet the spectrum of possible outcomes for a random process can be quite prescribed. For example, in rolling a pair of dice or flipping a coin there is no way to know the value of a specific outcome a priori but the likelihood of any given roll or flip is known exactly. Moreover, flipping a coin or rolling a die 1 million times will generate predictable distributions of outcomes. This seemingly simple condition has profound implications to assessing risk under certain conditions.

A sufficiently large number of randomly occurring events will ultimately yield a normal distribution of outcomes that are dispersed about a mean or average value. The normal distribution density function $g(x)$ for a random variable x is given by Eq. (1.2)[7]:

$$g(x) = \frac{1}{\sqrt{2\pi}} e^{-(1/2)x^2} \tag{1.2}$$

Crucially, the likelihood that a specific number of events have a certain value can be determined from this distribution. For example, by definition 68.2% of the outcomes of a normal distribution will occur within one standard deviation of the mean. The standard deviation specifies the dispersion or uncertainty about the mean of a distribution. The magnitude of the dispersion, which corresponds to some number of standard deviations about the mean, is the same for any normal distribution.

Fig. 1.6 is a graphic depicting a normal probability distribution of some population of events or things, and the fraction of the population corresponding to one, two, and three

[7]To arrive at a specific value for x requires integrating the density function from minus infinity to that number.

standard deviations about the mean, μ.[8] Note that these fractions and corresponding standard deviations are the same for any normally distributed random variable.

It is readily apparent from Fig. 1.6 that 68.2% of the distribution of values falls within one standard deviation of the mean, 95.4% of the distribution of values falls within two standard deviations of the mean, and 99.8% of the population distribution falls within three standard deviations of the mean.

The importance of the normal distribution to every field of science and to security risk management in particular cannot be overstated. Reference is made to this distribution throughout this text, and its properties will be invoked when discussing the Probability of Protection method in Chapter 13.

The normal distribution is not the only distribution that applies to random variables. The Poisson distribution is another worthy of mention. It is used to model discrete, randomly occurring events and is predicated on three assumptions[9]:

1. The probability of one event occurring in a time interval, $\Delta\tau$, is proportional to $\Delta\tau$ when $\Delta\tau$ is very small.
2. The probability that more than one event occurs in the time interval $\Delta\tau$ is negligible when $\Delta\tau$ is very small.
3. The number of events that occur in one time interval is independent of the number of events that occur in another nonoverlapping time interval.

One useful property of the Poisson distribution is that the mean and the standard deviation are the same value. The probability density function for Poisson distributions is given by Eq. (1.3) where λ is the arrival rate, k is a specific number of events, e is Euler's number, that is, 2.71828…, and $k!$ is "k factorial" or $k \times (k-1) \times (k-2) \times (k-3)$, etc.:

$$P(k;\lambda) = \frac{(\lambda)^k e^{-\lambda}}{k!} \tag{1.3}$$

The expression yields the probability that k events occurs in a given time interval assuming a constant event arrival rate λ.

The normal and Poisson distributions are related. In fact, for a sufficiently large number of events, a Poisson distribution morphs into a normal distribution. Fig. 1.7 is a graphic depicting the Poisson distribution for different mean values.[10]

Poisson distributions involve randomly occurring, discrete events that adhere to the three assumptions noted above. Examples of such processes include radioactive decay and photon counting. Chapter 9 discusses photon counting in the context of an optical attack on information assets.

A simple example of the use of the Poisson distribution is illustrative. Suppose one was developing a crude packet detector for a 1-Gb/s Ethernet interface in order to detect nascent denial-of-service attacks. Let us assume that packet arrival at the Ethernet interface is a random variable and obeys Poisson statistics as dictated by the three conditions noted earlier.

[8]http://www.srh.noaa.gov/bro/?n=2009event_hottestjuly.

[9]https://www.cis.rit.edu/class/simg713/Lectures/Lecture713-07.pdf.

[10]http://earthquake.usgs.gov/learn/glossary/?term=Poisson%20distribution.

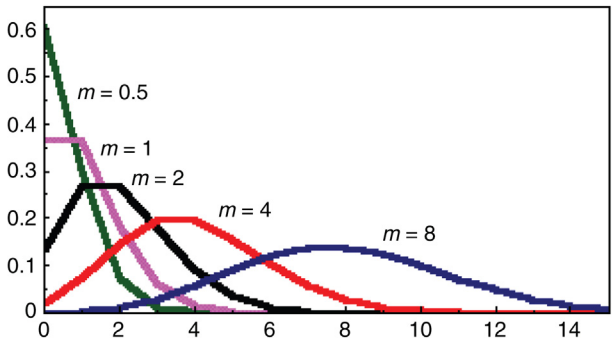

FIGURE 1.7 **Poisson distributions with different means.**

Suppose further that our detector counts packets in a given time interval t where a mean packet arrival rate λ is known for that link. However, the packet-forwarding rate for this size Ethernet link is assumed to be 1 packet/s (p/s).[11] If the detector is sized to expect the mean packet arrival rate, what is the probability the detector will see precisely 10 packets in a 5-s interval?

The Poisson density function is given by (1.3), and the calculation yields a probability of about 2%. The probability the detector would see less than or equal to 10 packets, that is, the cumulative risk, is about 99%. Therefore, if the detector registers significantly more packets in a 5-s time interval, it should register an alert since this behavior is not representative of normal system conditions.

The standard deviation of both normal and Poisson distributions is proportional to the square root of the sample size, N. What is the implication to estimating N to a specific level of precision?

The standard deviation, σ, is proportional to $\sqrt{N}$, where N is the sample size. Suppose it is mandated that the standard deviation of a distribution be one-tenth the sample size or $N/10$. This condition coupled with the dispersion about the mean associated with normal distributions determines the sample size that is required to achieve this level of precision.

Specifically,

$$\sigma = \sqrt{N} = \frac{N}{10}$$

Therefore,

$$\sqrt{N} = \frac{N}{10} \quad \text{or} \quad N = 100$$

[11]The actual packet-forwarding rate for a 1-Gb/s Ethernet link is between 81,274 and 1,488,096 p/s because the number of bytes/packet varies. However, the numbers have been scaled way back to facilitate a simple calculation, which hopefully will not diminish its instructive value.

Suppose the standard deviation is now specified such that it can be no greater than one-hundredth the value of N. This specification would require that $N = 100^2$ or 10,000.

Now in order to achieve 10 times that precision or $N/1000$, N must equal $(1000)^2$ or 1,000,000. From these examples one can generalize that the required precision scales as the square of the sample size.

So statements on the certainty of a specific outcome can be made for distributions of random variables. More specificity regarding the probability that a particular outcome is within a certain distance from the mean can be achieved for distributions with larger sample sizes. *However, such statements are valid only for distributions of random variables.*

To be clear, information security events do not occur randomly in the same way that some physical processes do such as in the production of photons in radioactive decay. To complicate matters, security incidents are relatively rare, and the conditions associated with each incident will vary. This is what makes it difficult to establish a probability distribution of security incidents versus loss. Other industries have created such distributions and one in particular is worth analyzing in some detail.

Investment banks are required to calculate and report their so-called Value at Risk (VaR), which drives their requirement to maintain capital reserves. The VaR is a broadly adopted metric in the banking industry. It corresponds to the probability of loss on a specific portfolio of financial exposures.

For a given portfolio, the time horizon and the p(VaR) is defined as a threshold loss value, such that the probability that the loss on the portfolio over the given time horizon exceeds this value is p. This model assumes so-called mark-to-market pricing, and no trading in the portfolio [5].

For example, if a portfolio of stocks has a 1-day, 10% VaR of $1 million, there is a 0.10 probability that the portfolio will fall in value by more than $1 million over a 1-day period assuming no trading. Alternatively, a loss of $1 million or more on this portfolio is expected on 1 out of 10 days.

Conversely, trading losses over a 1-day period would be expected to be less than $1 million 90% of the time. Clearly financial institutions are incentivized to reduce the amount of money held in reserve since more funds would be available for generating revenue. So the method used to calculate VaR has implications to the bottom line.

For financial institutions the capital reserve calculation is subsumed under the general heading of "Operational Risk." Fig. 1.8 shows a loss distribution curve specifying the portions of the curve corresponding to the Expected Loss, Unexpected Loss, the Extreme Loss, and VaR [6].

Establishing such a curve is the Holy Grail for security risk assessments. Unfortunately it is often elusive in security contexts because precise statements on incident probability versus loss are not so easy to determine. In addition, the value of information is not like the price of bananas, coffee, gold, machinery, equities, fixed income products, etc. Even financial institutions struggle with VaR, because large if infrequent losses (i.e., "tail events") are difficult to model.

Some types of losses are difficult to quantify in security contexts. As a trivial example, the electronic compromise of a 911 (999 in Britain) calling system might result in the loss of 10 lives because of slow response times by emergency services. Where would the magnitude of that loss appear on the VaR curve?

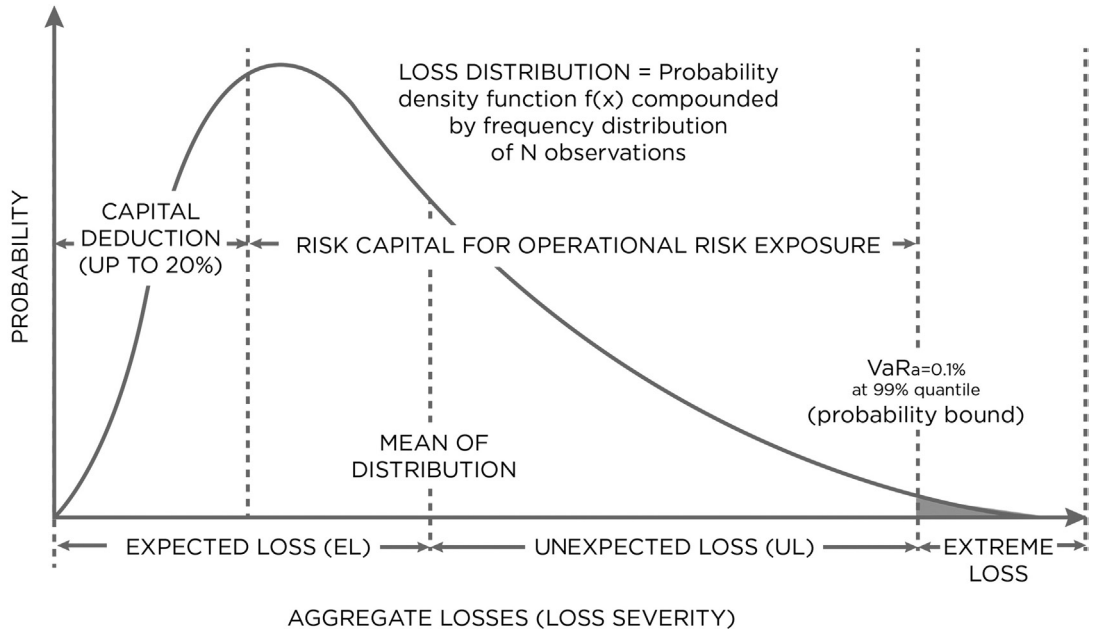

FIGURE 1.8 **Loss distribution approach (LDA) for operational risk under the New Basel Capital Accord.**

In addition, the same type of attack can result in wildly dissimilar outcomes for a given organization based on slightly different controls or network configurations. Attacks might result in only embarrassment because the objective of the attacker was not to steal information, yet large losses could have been easily sustained if the attackers were so inclined.

The applicability of VAR to real-world security scenarios is questionable for reasons previously cited and should be subject to significant scrutiny before its adoption [7].

The earlier discussion is not meant to imply that meaningful assertions about the *relative* likelihood of a future security incident are not possible. For example, the likelihood of an attempted denial-of-service attack against a high-profile institution such as a bank would likely be higher than for a relatively unknown Mom and Pop store.

In general, information security attacks against specific organizations are more likely because of the existence of specific risk factors such as their Internet presence, the value of the information they store, their political profile, and affiliations with specific individuals or entities. One can therefore safely say the *potential* for a future security incident is increased relative to other organizations based on one or more of these risk factors.

One might argue that there is an increased potential for attack based on Internet activity alone. In fact, the very existence of network connectivity might be considered a risk factor for information compromise since sharing information is the very objective of a network. Moreover, anything that facilitates sharing increases the vulnerability to compromise.

One should not interpret the foregoing discussion as an excuse to ignore the likelihood component of risk in conducting information security risk assessments. If one managed to obtain data from a statistically significant sample of organizations relative to parameters such

as network configurations, infrastructure, and the nature of existing security controls during successful attacks, one could make meaningful statements on the vulnerability as a function of scenario-specific parameters.

For example, a statement such as "60% of all successful computer system intrusions occurred when relevant account passwords were 10 characters or less in length" might be useful in developing a security strategy. This statistic would not necessarily demonstrate cause and effect, but it represents a strong correlation, and therefore should give a security risk manager plenty of food for thought.

It is true that certain attackers are known to exploit vulnerabilities merely because they are able to do so and do not necessarily have a reason for targeting a specific entity. However, in general, various risk factors such as those noted earlier increase an organization's attractiveness as a target and should be factored into the overall assessment of the magnitude of information security risk.

SUMMARY

Threats are entities, conditions, or phenomena that cause harm, damage, and/or loss. Risk is a fundamental characteristic of all threats, and is what makes a threat "threatening" to each organization. It therefore provides the context that enables prioritization of remediation efforts.

Risk has three components: impact, vulnerability, and likelihood. Impact refers to the importance of a threat incident, vulnerability is the magnitude of potential loss or the exposure as a result of a threat incident, and likelihood is the probability of a future threat incident occurrence. Successfully evaluating the three components of risk enables the prioritization of risk mitigation efforts.

Importantly, risk factors are features that enhance one or more components of risk with respect to a specific threat. Effective risk assessments must focus on identifying the spectrum of risk factors and identifying mitigation measures that address each of them with limits dictated by the assessed risk and available resources.

Varied conditions and the lack of controlled experiments make predictions of future information security threat incidents difficult. A risk-based and therefore contextual information security policy with accompanying IT standards should provide the basis for rigorous security risk assessments.

References

[1] Young C. The science and technology of counterterrorism; measuring physical and electronic security risk. Waltham, MA: Butterworth-Heinemann; 2014.
[2] Cost of developing a new drug. Tufts Center for the Study of Drug Development, Tufts School of Medicine, November 18, 2014. <http://csdd.tufts.edu/files/uploads/Tufts_CSDD_briefing_on_RD_cost_study_-_Nov_18,_2014.pdf>.
[3] Young C. Metrics and methods for security risk management. Boston: Syngress; 2010. p. 45–75.
[4] NIST framework for improving critical infrastructure cybersecurity, version 1.0; February 12, 2014.
[5] Jorion P. Value at risk: the new benchmark for managing financial risk. 3rd ed. New York: McGraw-Hill; 2006.
[6] Jobst A. The sting is still in the tail, but the poison depends on the dose. IMF working paper. <https://www.imf.org/external/pubs/ft/wp/2007/wp07239.pdf>.
[7] Jaisingh J, Rees J. Value-at-risk: a methodology for information security risk assessment. West Lafayette, IN: Purdue University. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.108.5140&rep=rep1&type=pdf>.