CYBERSECURITY FOUNDATIONS
AN INTERDISCIPLINARY INTRODUCTION

CHAPTER 1: Risk Management for Cybersecurity

CHAPTER OUTLINE

- Overview of Risk Management
- Introduction
- Risk Management Frameworks
- The Process of Risk Management
- Risk Framing
- Risk Assessment and the Risk Formula
- Threat Assessment
- Vulnerability Assessment
- Consequence Assessment
- Risk Determination
- Risk Response
- Risk Monitoring
- Conclusion
- Key Questions

1 Risk Management for Cybersecurity

Chapter In Focus

- **The theory, process, and practice of risk management for organizations and individuals.**
- **The assumptions, constraints, tolerance level, priorities, and trade-offs involved in risk framing.**
- **The three assessments—threat assessment, vulnerability assessment, and consequence assessment—that inform the process of assessing overall levels of risk.**
- **Qualitative and quantitative approaches to risk determination, including the probabilistic risk assessment (PRA).**
- **The possible variations of risk response: risk acceptance, risk avoidance, risk mitigation, and risk transfer.**
- **The steps of the ongoing risk monitoring process: compliance, effectiveness, and identification of changes.**

Overview of Risk Management

Risk management, the process by which risk specialists develop and implement a continuous and systematic plan for containing risk, is the practice at the heart of cybersecurity. Risk managers use words like "manage" and "contain" to describe their work because in the cyber realm, as in the physical world, risk can be managed but never completely eliminated.

Cyberspace is a risk-laden virtual environment. The possibilities for disaster and data theft are infinite in cyberspace: systems fail, hackers are omnipresent and always developing new methods of attack, and physical damage can destroy servers and databases. Each time you use a device with an Internet connection, you are putting your personal information at risk. Therefore, to secure operations in cyberspace, risk managers and organizations must think through, select, and implement a customized risk management plan and continuously monitor the plan's successes and failures.

Risk management strategies for the physical world, including plans for national security emergencies, have influenced risk management strategies for cyberspace operations. Before the terrorist attacks of September 11, 2001, national security experts had not considered the possibility that terrorists might fly planes into skyscrapers. Because this kind of attack was so improbable as to be inconceivable, national security risk managers never developed a comprehensive risk management strategy to prepare emergency responders, civilians, and national leadership for the possibility of an airplane attack on civilians. But 9/11 revealed to the United States and the world the dangers of disregarding the risk of improbable events. In the post-9/11 world, both private and public

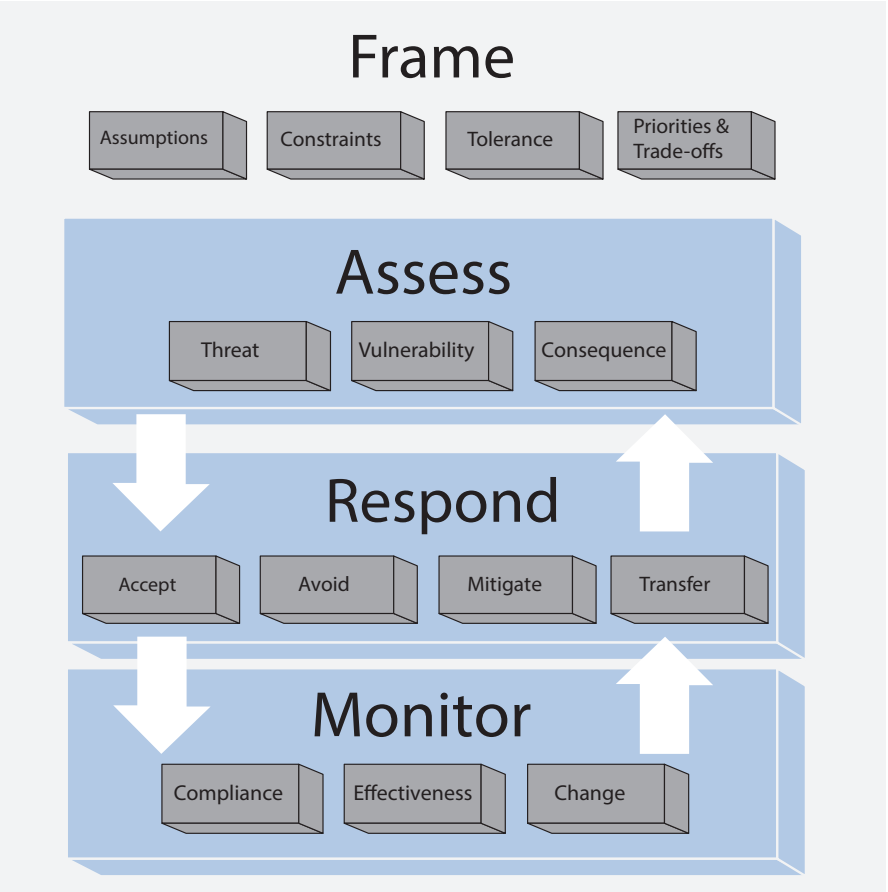

Figure 1-1: Risk framing process.[1]

sector risk management strategists now include calculations for highly improbable events into their planning and strategies. Moreover, the devastation of 9/11 has impressed upon the cybersecurity community the need to prevent a "cyber 9/11."

Because risk can never be eliminated, risk management is a field in which innovations and adjustments are always possible and often necessary. Each recent national catastrophe the U.S. has faced, from Hurricane Katrina to the financial meltdown of 2008, has influenced risk management models at the highest levels of both the federal government and the private sector. Cyber risk management models now frequently incorporate elements from both government and private sector strategies. Nevertheless, because of the complexities and unique challenges facing individuals, organizations, businesses, and governments operating in cyberspace, risk managers for cyber operations must develop highly specific plans to contain cyber risk.

Introduction

Risk management is a process that formalizes the steps of identifying the most critical assets to an individual, organization, or company; assessing the risk; and determining the best method to prevent harm to these assets. Risk is the likelihood that a given threat will exploit a particular vulnerability and the resulting

consequence of that negative incident occurring. Risk management permits cybersecurity specialists to weigh technical and financial costs of security measures that support the organization's function.

Some kinds of risk management are intuitive: when you approach a street crossing, you know that you face the risk of being hit by a car. You protect yourself by looking both ways and assessing the situation before deciding how to proceed. Similarly, if you're deciding whether to come down from a tree by jumping or using a ladder, you consider the potential consequences of each choice and minimize the possibility of bodily harm by opting for the ladder.

In neither of these intuitive examples is the process formalized. Both examples lack (and do not require) a thoughtful and systematic approach to the process of protecting the most critical asset at stake, your physical health.

In contrast, risk management as a professional practice follows a codified process. A schematic representation of this process is shown in Figure 1-1. Note that the process is not strictly linear. Each of the four major elements—"Frame," "Assess," "Respond," and "Monitor"—is interlinked with the others, and the process is continuously changing and reacting to circumstances in real time.

Risk Management Frameworks

The **National Institute for Standards and Technology (NIST)** has developed various risk management frameworks based on the intended system or organization. NIST is a federal agency that works with technology companies to create standards and procedures. These standards and procedures establish baselines for the deployment of products or services across an industry. All industries work with NIST to develop industry-wide standards for the products and services they deliver.

Cyber Fact

NIST Special Publication 800-30 defines an "IT system" as a "general support system (e.g., mainframe computer, midrange computer, local area network, agencywide backbone) or a major application that can run on a general support system and whose use of information resources satisfies a specific set of user requirements."

Although this chapter loosely follows the NIST model for risk management, it is important to note that alternative frameworks exist. Within the federal government, in fact, each department and agency has its own methodology, whether formalized or not. One important alternative to NIST is the National Infrastructure Protection Plan (NIPP), established by the Department of Homeland Security (DHS). DHS has also adopted a **Risk Lexicon** that we will refer to frequently in the course of this chapter.[2]

The Process of Risk Management

This chapter will explain the four major steps of the risk management process: framing, assessing (including determining), responding, and monitoring risk. We will dissect each step using two non-cyber examples—a group of friends driving to a party or nightclub on New Year's Eve and a homeowner dealing with the possibility of burglary—as well as a cybersecurity example—an organized crime group that targets e-commerce sites, predominately commercial retail-

Cyber Fact

An **organized crime group** is a domestic or international non-state actor with a political, social, or economic ideology. The group impresses its ideology on its targets through illegal activities and/or violence. Lulz Security (LulzSec) and its affiliate, Anonymous, are infamous organized crime groups that have been implicated in major cyber crimes in recent years. Other organized crime groups are purely interested in financial gains or other types of resources.

ers, with the intent of disrupting commerce and stealing customer information. The friends on New Year's Eve face the risks associated with drunk and reckless drivers, in addition to the everyday risks of traffic accidents and automotive malfunction; these risks represent unintentional threats, because drivers do not ordinarily intend to harm one another. The family intent on securing its house and property against burglars, in contrast, faces intentional threats, because burglars don't burgle by accident.

As we'll see, cyber risk managers must consider both intentional and unintentional threats at each step in their work process. Based on our hypothetical organized crime group's sophisticated hacking tools and vast financial resources, the group is able to conduct a large-scale and devastating attack against an online retailer. For the purposes of our discussion, this cyber example temporarily suspends the concept of regulations or legal requirements with which the company must comply (to learn more about these requirements, see Chapter 5).

Figure 1-2 shows a systematic breakdown of the cyber example with specific examples of activities at each phase. We suggest that you consider Figure 1-2 now and then return to it after you have worked through the chapter.

Risk Framing

The first step in the cyber risk management process is **risk framing**. Risk framing is the process of examining and evaluating the "big picture" risk environment in which a company or organization operates. Risk framing establishes the context for making **risk-based decisions**. Risk-based decisions, according to the DHS Risk Lexicon, are "determination[s] of a course of action predicated primarily on the assessment of risk and the expected impact of that course of action on that risk." Organizations and individuals routinely make risk-based decisions that affect investments and operations. For example, when a company allocates resources to managing risk, this decision means that resources are taken away from other program areas. Risk framing establishes the context in which the risk must be managed and establishes the risk-based restrictions around organizational decisions. The purpose of risk framing is to produce a risk management strategy.

Risk framing examines the assumptions, constraints, tolerance level, and priorities and trade-offs associated with risk. Figure 1-3 considers these elements of risk framing as they pertain to our example of an organized crime group that sends malicious code to a large company in an attempt to disrupt commerce and obtain personal customer information.

Assumptions

Assumptions are the reasonable expectations of actions, tools, and policies that may already be in place to protect critical assets, or reasonable expectations of the risk faced. The friends on New Year's Eve make the assumption that it is riskier to drive on New Year's Eve than on other nights. The homeowner and his family as-

The Scenario:	An organized crime group sends malicious code in an email to a popular U.S. company known for the high volume of traffic on its website. The organized crime group wants to crash the company's website, which will disrupt commerce.

1. Frame Risk

Assumptions	Anti-virus software has been updated and installed on all company computers. This software will prevent malicious intrusions.
Constraints	Not all of the company's financial resources and personnel can be devoted to securing the website.
Tolerance	A major U.S. company that relies on web sales has a low tolerance for disruption to its website.
Priorities & Trade-offs	Protecting the website from being disrupted is a high priority for the company. Therefore, the company must take action to avoid the disruption of online business. As a result of the funds dedicated to website security, another office within the company will face budget cuts.

2. Assess Risk

Threat	An organized crime group emailing malicious code that could seriously disrupt the company's web functioning and sales.
Vulnerability	An individual within the company, unaware of the repercussions, opens an email containing malicious code.
Consequence	Financial loss, disruption of service, and harm to the company's reputation.

3. Risk Response

Accept	Receive emails from unknown users.
Avoid	Stop using the company's computers.
Mitigate	Protect the server with a firewall.
Transfer	Purchase cyber insurance.

4. Monitor Risk

Compliance	Verify the implementation of a company-wide firewall.
Effectiveness	Continually test and ensure that the firewall is working.
Identify Changes	Identify new threats and vulnerabilities, critical assets, and consequences as the risk environment changes and evolves.

Figure 1-2: Risk management outline and examples.

1. Frame Risk	
Assumptions	Anti-virus software has been updated and installed on all company computers. This software will prevent malicious intrusions.
Constraints	Not all of the company's financial resources and personnel can be devoted to securing the website.
Tolerance	A major U.S. company that relies on web sales has a low tolerance for disruption to its website.
Priorities & Trade-offs	Protecting the website from being disrupted is a high priority for the company. Therefore, the company must take action to avoid the disruption of online business. As a result of the funds dedicated to website security, another office within the company will face budget cuts.

Figure 1-3: Risk framing elements and examples.

sume that their home is a target for burglars. In our cyber example, there may be an assumption that antivirus software is installed and working correctly. For government employees, another assumption might be the presence of a **Trusted Internet Connection (TIC)** (see "Cyber Fact" on the following page). However reasonable these security assumptions may seem, they still must be checked, tested, and confirmed before an organization (including, in this case, a group of friends, a family, or a company) can develop and formalize its strongest possible risk management plan.

Constraints

Risk constraints are the factors that inhibit the execution of a 100-percent secure risk management plan. The most common risk constraint in many scenarios is the finite nature of financial resources. Organizations and individuals operate on restricted budgets, and they can't afford to spend all of their resources to protect a critical asset. For this reason, organizations and individuals must distribute resources strategically and intelligently to ensure that all of their critical assets have some measure of appropriate protection from harm.

Cyber Fact

The Office of Management and Budget mandated the use of the **TIC** in order to decrease Internet access points on government networks and to verify that all access to these networks is routed though designated TIC Access Providers. The TIC initiative is intended to "optimize and standardize individual external network connections currently in use by federal agencies, including connections to the Internet."[3]

In the New Year's Eve example, a constraint may be that the group of friends cannot drive an indestructible tank through a modern city in order to reach its destination. One of the family's constraints is that it only has a certain amount of money to spend on a home security system. In the cyber example, the company cannot shift its entire budget to reduce the threat of an organized cyber attack; the company, like the family, is financially constrained. Another possible constraint for the company is the need to maintain a user-friendly website. Online retailers must strike a balance between website security and website accessibility; customers may eschew a website that has too many security protocols in place because additional security measures may feel cumbersome and delay transaction time.

Tolerance

Risk tolerance is the degree to which an organization can handle or incur a specific harm. A car has a lower harm tolerance in the event of a collision than a tank does. A home without an alert system has a lower risk tolerance than a home with an alert system. The risk tolerance level of a situation or organization depends upon both the situation's total vulnerability and the threat's potential to inflict harm. We will discuss the critical risk management concepts of threat and vulnerability later in this chapter.

A measurement of risk tolerance offers information about whether or not a potential risk is acceptable to an individual, organization, or entity. A low risk tolerance level indicates an unacceptable risk. According to the DHS Risk Lexicon, unacceptable risk is the "level of risk at which, given costs and benefits associated with further risk reduction measures, action is deemed to be warranted at a given point in time." Unacceptable risks must be addressed with action and resources. In our cyber example, the company's tolerance for the risk of receiving a malicious code embedded in an email is greater when the code comes from a lone hacker rather than when it comes from an organized crime group. The company can tolerate the risk posed by a lone hacker, because this hacker probably does not possess the capability to crash the company's website singlehandedly; the lone hacker probably also has fewer resources than the organized crime group. The company cannot tolerate a sophisticated hacking ring, but it can probably tolerate the activities of a lone hacker. Therefore, if the company has reason to believe that its network is the target of a cyber attack by an organized crime group, it must take action and commit resources to contain this unacceptable risk.

Priorities and Trade-offs

As the final step in a comprehensive risk framing process, an organization or individual identifies **priorities and trade-offs** that must be negotiated between all stakeholders in the scenario. For the individuals in our New Year's Eve example, priorities include safe and timely arrival at their destination as well as minimization of transportation costs. If the group designates a non-drinking driver, a trade-off is that the driver must remain sober and alert throughout the evening. The family decides that a priority is to secure its house from a burglar; as a result, they might have to make the trade-off between devoting money toward a new alert system and spending it on a vacation.

The company in our cyber example must make a similar financial trade-off if it hopes to contain the major risk represented by the organized crime group's potential to launch a cyber attack.

Indeed, the company facing the threat of malicious code crashing its website must prioritize the containment of this specific risk. Understandably, the company ranks these kinds of massively disruptive cyber threats as a primary concern, so it dedicates extra funds to staff, programs, and technology to protect its network. The trade-off for the company is that prioritizing these cybersecurity concerns reduces the funds available to other activities, such as employee training, marketing, or business development. Prioritizing security concerns against other business or management issues can be a difficult decision for executives to make. Ultimately, managers must make trade-offs based on the assumptions, constraints, level of tolerance, and priorities identified during the risk framing process.

Risk framing is the first step in creating a strategy to reduce risk. Thorough risk framing also provides a long-term strategic view of an organization's decision-making process. The strategy that an organization generates from the risk framing discussion outlines the challenges and mechanisms involved in protecting its most critical assets from harm.

Risk Assessment and the Risk Formula

The second phase in the cyber risk management process is **risk assessment**. The DHS Risk Lexicon defines risk assessment as a "product or process which collects information and assigns values to risks for the purpose of informing priorities, developing or comparing courses of action, and informing decision making." The purpose of a risk assessment is to reveal the factors that, taken together, constitute a full picture of the risk that an organization faces. The full risk assessment consists of three distinct elements: threat, vulnerability, and consequence assessment. If a threat and/or vulnerability exists, then risk exists.

It is often easier to consider the three assessments together as one risk assessment. In many risk assessments, the three components are virtually inseparable from each other. Threat assessments and vulnerability assessments, in particular, can be difficult to separate. Nevertheless, a careful consideration of threat, vulnerability, and consequence as individual components of a risk assessment will yield a stronger, more nuanced assessment.

The Risk Equation

The generic formula for the summation of risk is shown in Figure 1-4. This equation cannot be used to calculate the likelihood of a specific harm occurring. Later in this chapter, we will consider equations that can be employed to produce hard calculations of specific risks.

Risk = Threat × Vulnerability × Consequence

Figure 1-4: The risk equation.

To understand the risk that a company or an individual faces in any given situation, we must identify and measure threats, vulnerabilities, and consequences inherent in the situation. None of these three aspects of the risk assessment can be fully effective without the other two. All three of these words – threat, vulnerability, and consequence – have casual meanings for everyday use that are related to, but not exactly the same as, their technical meaning for the process of assessing risk. For risk assessment purposes, *threats* are the agents that cause harm to an individual or organization. *Vulnerabilities* are identifiable weaknesses in processes, personnel, networks, or other technologies. The *consequence* is the specific result that occurs if a threat exploits a vulnerability and causes damage or harm to an individual or organization. Consider the following:

- The threat of a drunk driver on New Year's Eve; the threat of a burglar invading a home; the threat of an organized crime group bringing down a company's website.
- The vulnerability inherent in driving a car; the vulnerability of living in a particular home; the vulnerability of the IT network architecture at an online company.
- The consequence of a car accident on New Year's Eve; the consequence of a burglar entering a family's home; the consequence for the online company of an attack on its computer network and company website.

Once we conduct separate threat, vulnerability, and consequence assessments for these situations, we will possess a unified and comprehensive view of the risk facing the group of friends driving on New Year's Eve, the family trying to protect its home from a burglar, and the online company facing cyber threats from an organized crime group. After conducting a risk assessment, we will have the indicators necessary to determine the severity and magnitude of a risk, as well as the general (non-specific) likelihood of a risk occurring.

Threat Assessment

The DHS Risk Lexicon defines a **threat** as a "natural or man-made occurrence, individual, entity, or action that has or indicates the potential to harm life, information, operations, the environment, and/or property." In other words, threats are the agents that cause harm to an organization's processes, systems, personnel, hardware, software, or physical location. Threats cause harm by exploiting vulnerabilities.

According to DHS, a **threat assessment** is a "product or process of identifying or evaluating entities, actions, or occurrences…that have or indicate the potential to harm life, information, operations, and/or property." The purpose of the threat assessment is to identify the intention (target or goal), capability (power to commit harm), and lethality (level of harm) of a threat.

As we discussed earlier, a threat may or may not be intentional. Risk managers must consider both *intentional* and *unintentional* threats at this stage in the risk assessment. Types of intentional threats include physical or cyber attacks. The burglar entering the family's home is an intentional threat. In our cyber example, the crime group's intention may be to gain access to customer information. The group could use this information to open new credit card or bank accounts, or sell the information to third parties. Other cyber crime groups simply want to embarrass a company or organization, or make an ideological point by disrupting a business or government's ability to function.

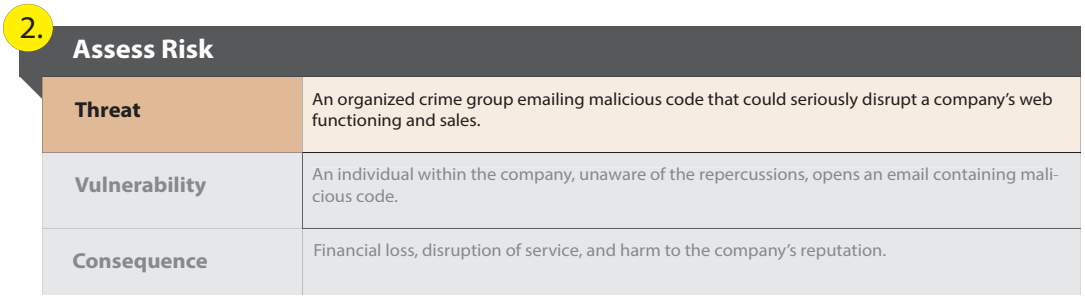

Figure 1-5: Risk assessment elements and examples.

Unintentional threats include human errors of omission or commission and natural or man-made disasters. A drunk driver on New Year's Eve may not intend to cause a fatal accident, but this lack of intention does not make his presence on the road any less dangerous. The individuals in the New Year's Eve scenario face many unintentional threats: the specific elevated threat of a drunk driver on New Year's Eve, in addition to the threats inherently associated with driving a car at any time (malfunctioning parts, distracted drivers, animals in the road, etc.).

Professional risk managers are charged with breaking down and analyzing threats, vulnerabilities, and consequences as separate factors in a risk assessment. In the case of national security questions, this task involves sophisticated algorithms and may take months or years to complete to a satisfactory level. Even at a basic level, threats and vulnerabilities may be difficult to distinguish from one another. The driver of the group of friends on New Year's Eve may lose concentration at the wheel. While this may at first seem to create a new vulnerability, the driver has also become an unintentional threat to himself, to the car, to the other passengers in the car, and to other cars, passengers, and drivers on the road.

In the case of the organized crime group, the threat has a capability level (power to commit harm) that is determined by the hacker's level of skill and sophistication, expertise, resources, and quality of tools for access-

ing secured systems. The lethality (level of harm) to the company could take the form of immediate financial loss, short-term and long-term loss of customers, damage to the company's reputation, and the need to allocate additional resources for security and public relations. In this specific case, the lethality of both financial and non-financial losses would be highly detrimental to this company's business if its website were temporarily rendered inoperable, or hacked into in order to extract customer information. The exact figure of financial loss or non-financial consequence is calculated in the consequence assessment and risk determination step of the risk management process. We will discuss both processes later in the chapter.

In our cyber example, it is not difficult to identify many ongoing threats inherent when using a computer for critical company operations. To offer just one of the many possible examples, there is always the threat of a pipe breaking and water harming the company's server. Another unintentional threat might be that a programmer makes an error that introduces a new vulnerability into the system. A strong cyber threat assessment identifies these kinds of internal threats in addition to external threats such as the organized crime group. Figure 1-6 provides a breakdown of different types of threats and the possible consequences that they may result in.

Finally, a threat assessment may reveal threat shifting. **Threat shifting**, according to DHS, is the "response of adversaries to perceived countermeasures or obstructions, in which the adversaries change some characteristic of their intent to do harm in order to avoid or overcome the countermeasure or obstacle." In other words, threat shifting occurs when an intentional threat actor becomes aware that mitigations or controls are in place to thwart its activities, prompting the threat actor to alter its strategy to achieve its purpose. Cyber attacks are constantly evolving in new ways to exploit computer systems and gain access to sensitive information in storage and transmission (see Chapter 4 for more on common cyber attacks and exploits). Threat shifting is a major reason why risk management is a non-linear, ongoing process demanding constant monitoring, evaluation, and starting over.

Threat Assessment

Threat	Potential Targets	Capability	Objectives	Consequences
Nation-State	- Nuclear facility - Power grid - Financial markets	- Highest level capability - Vast funding/resources - Highly organized	Political goals	Catastrophic damage to national and economic security, if essential services are disrupted or there is loss of life.
Organized Crime	- Credit card numbers - Social Security numbers	- Higher level capability - Well-funded/organized	Money	Catastrophic or severe damage to the posture of a private company, if attack disrupts corporate services or results in loss of customer data.
Kiddie Hacker	- Social media account - Website	- Lower level capability - Limited resources	Prestige, fun	Minimal or insignificant damage if a personal website is temporarily defaced.

Figure 1-6: Threat assessment chart.

National Security Threats	**Information Warrior**	Reduce US decision space or strategic advantage, chaos, target damage
	National Intelligence	Information to gain political, military, or economic advantage
Shared Threats	**Terrorist**	Visibility, publicity, chaos, political change
	Industrial Espionage	Comparative advantage
	Organized Crime	Retribution, financial gain, institutional change
Local Threats	**Institutional Hacker**	Monetary gain, thrill, challenge, prestige
	Recreational Hacker	Thrill, challenge

Figure 1-7: The threat spectrum.[4]

Vulnerability Assessment

The purpose of the **vulnerability assessment** is to reveal weaknesses in facilities, personnel, systems, networks, technology, or processes. The vulnerability assessment specifically looks for particular weaknesses that correspond to specific threats.

According to DHS's Risk Lexicon, a **vulnerability** is a cyber or "physical feature or operational attribute that renders an entity, asset, system, network, or geographic area open to exploitation or susceptible to a given hazard."

Just as threats are omnipresent and inevitable in nearly every kind of scenario, vulnerability is also a universal condition. For instance, all devices connected to the Internet are vulnerable to some form of cyber intrusion, just as all buildings are vulnerable to physical intrusion or damage. Therefore, the vulnerability assessment goes hand-in-hand with the threat assessment.

Cyber Fact

DHS defines a **vulnerability assessment** as a "product or process of identifying physical features or operational attributes that render an entity, asset, system, network, or geographic area susceptible or exposed to hazards."

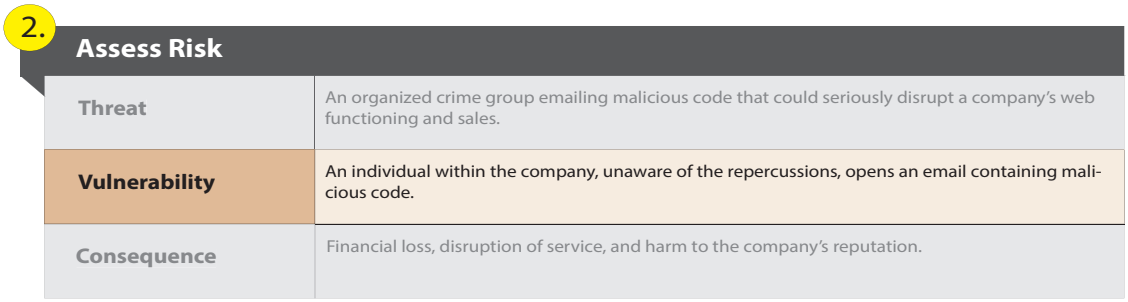

Figure 1-8: Vulnerability assessment example.

In our New Year's Eve example, the car's crushable exterior or lack of airbags may represent a vulnerability to the passengers. If the family's house lacks an alarm system, this lack may represent a crucial vulnerability in the event of a break-in. In our cyber example, a vulnerability may be a direct result of human error. Even though employees in the company may have received cybersecurity training and education, human susceptibility to error can never be completely eliminated. For example, an employee might open an email from an unknown source that runs a malicious code on his or her computer.

Additional vulnerabilities in the network architecture of a company may be the lack of an intrusion detection system or firewalls that are not strategically located to protect the organization's most critical assets. NIST has compiled a national database of information technology (IT) vulnerabilities that is helpful for all cybersecurity students and professionals to review.[5] Still, some IT vulnerabilities are more complicated than others, so NIST's list is not exhaustive.

In the event a vulnerability is exploited by a threat, the confidentiality, integrity, availability, or functions of critical assets may be compromised. Compared to the process of identifying threats, the process of identifying vulnerabilities may seem more challenging. Nevertheless, in many cases the threat's capability, lethality, and target or intention is specific enough to help narrow down the relevant vulnerabilities. Combined, the threat assessment and vulnerability assessment should yield a picture of the degree of potential harm to critical assets in the event that a vulnerability is exploited by a threat.

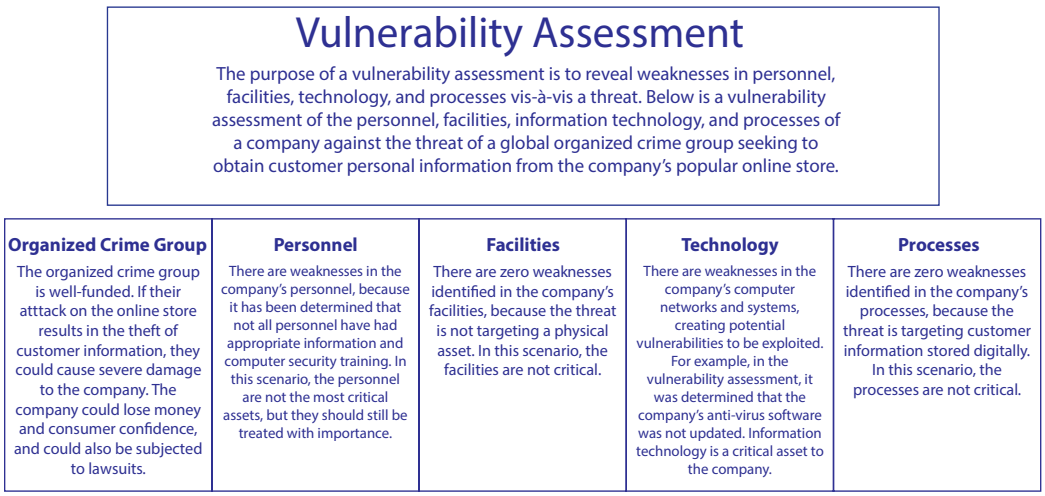

Figure 1-9: Vulnerability assessment diagram.

Cyber Fact

An **intrusion detection system (IDS)** is a monitoring device that identifies malicious activity in a network. A firewall protects a network by observing and detecting data, and then determining whether or not the data can pass through the network's firewall.

Consequence Assessment

The **consequence assessment** evaluates the potential impact on an organization in the event that a critical asset is exploited by a threat. **Consequence** is defined by DHS as the "effect of an event, incident, or occurrence." Consequence in risk assessment is determined by many factors, including the purpose and function of critical assets, the interdependencies between each critical asset and other assets, and how easy or difficult it would be to replace or repair each asset in the event of a breakdown or a lethal attack.

Both indirect consequences and direct consequences may ensue if a threat exploits a vulnerability. According to the DHS Risk Lexicon, an *indirect consequence* is an "effect that is not a direct consequence of an event, incident, or occurrence, but is caused by a direct consequence, subsequent cascading effects, and/or related decisions"; while a *direct or primary consequence* is an "effect that is an immediate result of an event, incident, or occurrence." For example, an indirect consequence to an online company attacked by an organized crime group is a damaged reputation; a direct consequence is a loss of money due to theft or lost sales during the time the website was down. The difference between the types of consequence may not always be this clear. Still, based on the scenario, the risk analyst will be able to define and separate the two types of consequence.

Before conducting the consequence assessment, we must answer the following question: "The threat will cause harm to which asset?" Determining and prioritizing an organization's critical assets is a key consideration for security policymakers and executives when they decide how to allocate limited resources in the face of conceivably unlimited threats. The value of each critical asset and the potential fall-out if any asset were to be harmed are the most crucial pieces of information that inform any consequence assessment.

Consequence is defined as the product of **criticality** and **impact**. *Criticality* refers to the importance of the critical asset and *impact* is the result of damage to the asset. A formula for the summation of consequence appears in Figure 1-10.

Consequence = Criticality × Impact

Figure 1-10: The consequence equation.

Let us examine ways of thinking about an asset's criticality or importance. One such method, highlighted by an official from the Government Accountability Office in an October 2001 statement before the House Subcommittee on National Security, Veterans Affairs, and International Relations, assigns the following labels to categories of criticality (in descending order of criticality): "catastrophic," "critical," "marginal," or "negligible."[6] These labels can then help decision-makers rank the potential criticality of the loss of or harm to any given asset.

Let us return to our two non-cyber examples to illustrate a consequence assessment. The harm to the individuals driving on New Year's Eve if a vulnerability (e.g., the car) is exploited by a threat (e.g., an intoxicated driver) is bodily damage to the passengers and damage to the car. The harm may be "marginal" if only the car is damaged, but the harm would be "catastrophic" if an individual is injured or killed in such a collision. If a burglar invades a family's home, the consequence may be defined as the value of the goods stolen. The harm is "marginal" if the items are replaceable within the family's budget, but it would be "catastrophic" if a family member is injured or killed in a fight with the burglar.

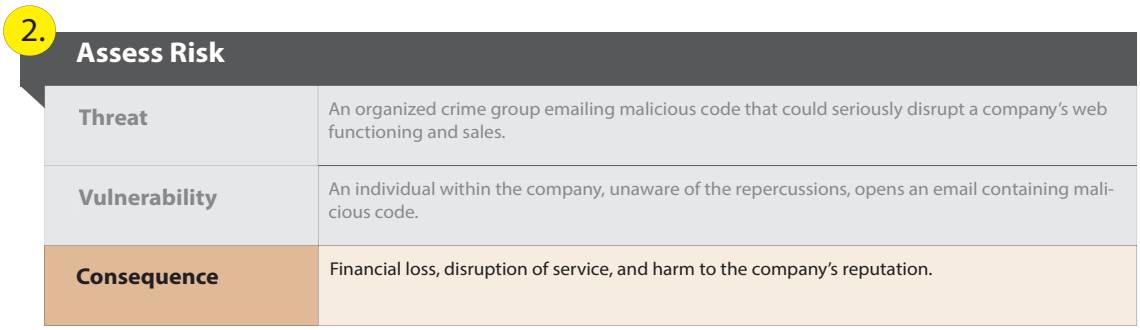

Figure 1-11: Consequence assessment example.

In our cyber example, the consequence of an employee opening malicious code that leads to the leak of customer information may be only "marginal" if fewer than ten customers are affected. The consequence may be "catastrophic" if thousands of customers have their information stolen and exploited; the indirect consequence of this scenario might be that the online retailer must pay damages to their customers and incur negative media attention.

Let us consider one more example: The economic and national security consequence and impact of a hacker causing catastrophic damage to New York City's power grid is greater than the consequence of the same hacker disrupting the power grid of a small town in North Dakota. In North Dakota, the consequence of such a disruption would be "marginal," whereas in New York City the consequence would be "catastrophic." The reason for these different levels of severity is that both the impact (more people affected) and the criticality (more important centers of industry) in New York City are greater than in any town in North Dakota. Because the consequence formula is criticality multiplied by impact, the consequence of an attack that damages or disrupts New York City's power grid will almost always be greater than the consequence of an attack on any other power grid in the United States.

These labels of criticality are subject to the particularities of each risk scenario, because each risk scenario redefines the consequence metric. For example, consequence may be expressed in terms of the hours a business's website is down, the legislation enacted as a result of a data breach, the financial loss to individuals after a hacking incident, or the number of people harmed in the event.

The combined threat, vulnerability, and consequence assessment is a valuable tool for identifying the nuances and complexities of the risk facing an individual, organization, government, or nation. We next turn to the step of risk determination, in which the specific risk that the organization faces is identified and analyzed.

Cyber Fact

According to its official website, **the Government Accountability Office (GAO)** "investigates how the federal government spends taxpayer dollars" and "advise[s] Congress and the heads of executive agencies about ways to make government more efficient, ethical, equitable, and responsive."

Risk Determination

Risk determination is the step in the risk management process that follows the risk assessment. It is often the most in-depth and technical step in the entire risk management process. Many risk methodologies include risk determination as part of the risk assessment, but it can be beneficial to consider it as a separate process that focuses on the risk in greater depth than the assessment.

Nevertheless, a focused and effective risk determination is contingent upon a solid and thorough risk assessment. For example, by first determining the various threats and vulnerabilities an organization faces, as well as the intention, capability, and lethality of the threats, a risk management team can decide which threats and vulnerabilities to focus on, thereby devoting its energies to containing the most significant risks. A comprehensive risk assessment may alter the initial risk management strategy by tweaking the risk manager's original assumptions, constraints, priorities, and trade-offs from the risk framing process, or by drawing attention to the organization's tolerance for a specific kind of harm. After a threat assessment and a fresh understanding of the capabilities of a threat, an organization's first assumptions about the kind of risk it faces will likely change. The risk determination must acknowledge and build on the results of the risk assessment.

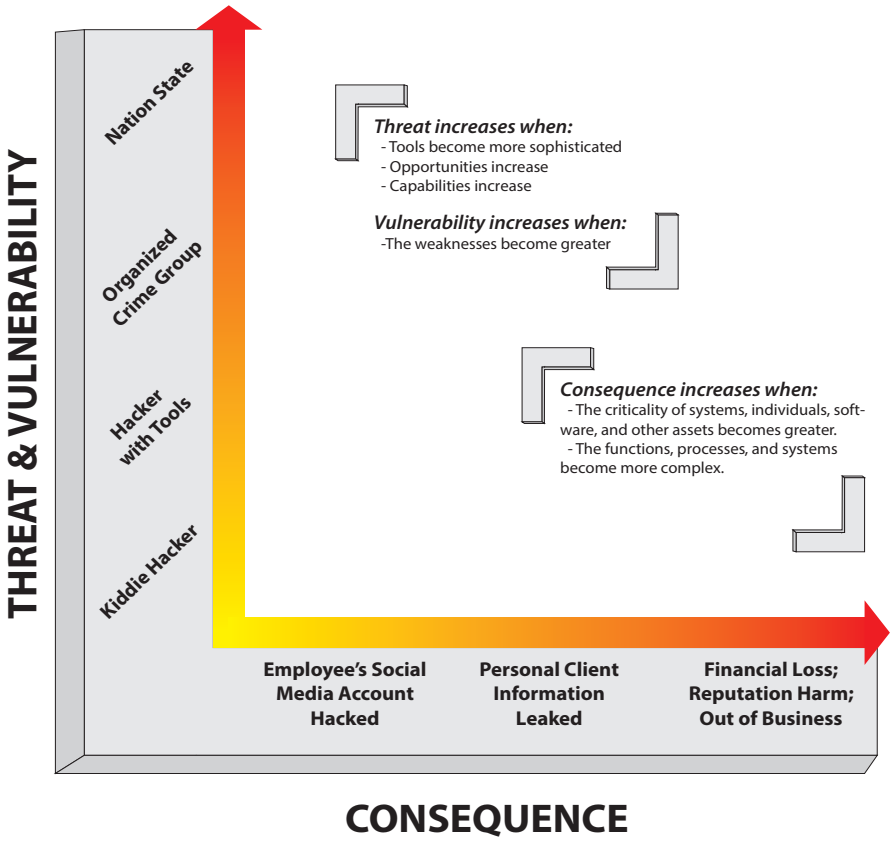

Figure 1-13: Consequence versus threat and vulnerability.

A risk analyst's first major insight into the level of risk an organization faces comes from examining threat and vulnerability versus consequence. Figure 1-12 (on page 36) illustrates different levels of risk. From a risk analyst's standpoint, the ideal quadrant for risk is Quadrant I; the most severe and undesirable level is Quadrant IV. The chart in Figure 1-12 will help guide our understanding of risk determination.

Because threat and vulnerability go hand in hand, when a risk analyst evaluates the level of risk an organization faces, threat and vulnerability are considered together. In Quadrant I, the combined threat and vulnerability level is low; therefore, the consequence level is also low. In Quadrant IV, the threat and vulnerability level is high; therefore, the consequence level is also high. If an organization's risk level is determined to be in Quadrant IV, then the risk is severe and the organization must address it immediately. When an organization's risk level falls in Quadrant II or Quadrant III, risk analysts and risk managers must decide how best to approach the risk. Risk analysts will evaluate the options for controls and mitigations available as part of a customized security program to address the specific risk.

Figure 1-13 provides insight into the progression of the level of risk. The horizontal axis describes the increase in consequence to the most critical assets. The vertical axis describes the increase to both the level of threat and the level of vulnerability. As a risk analyst thinks through this chart, he must consider what is happening along both axes simultaneously. The graph in Figure 1-13 is the visual representation of the equation we saw in Figure 1-4: *Risk = Threat x Vulnerability x Consequence*.

As you consider this graph, keep in mind that the terms of the vertical axis (threat and vulnerability) are not fixed. Nation states all possess different capabilities from one another, and some organized crime groups may have greater capabilities than some nation states. However, in general, nation states have the greatest capabilities (resources, people, and skills) to devote to a cyber attack. For this reason, many countries have recently devoted defense and civilian resources to educating their populations about cybersecurity issues, such as the issue of preventing cyber attacks on national networks, in particular.

Thus, the risk determination process provides a complete overview of the possible levels of risk and provides a guide to the appropriate risk response at each particular level. If the risk level is low and appears in Quadrant I, then the risk response plan does not need to be as robust as it would be if the risk level appears in Quadrants III or IV.

However, the likelihood of a specific risk occurring is still unknown. To determine the likelihood of a risk occurring, we rely on the total risk equation: the probability of an event occurring times its expected consequence value. A formula for the calculation of total risk appears in Figure 1-14.

Total Risk = Probability of an Incident Occurring × Expected Consequence

Figure 1-14: The total risk equation.

In order to calculate the precise mathematical likelihood of risk occurring, we require a numerical value for the consequence. As discussed, consequence is assigned a value, such as the number of hours a website is defaced or the amount of customer information leaked following an attack. The consequence value may be a combination of several factors. Figure 1-15 displays a graph representing risk likelihood.

The relationship illustrated by this graph is not necessarily linear; that is, the probability of an event occurring does not have to increase as the consequence increases. However, the level of risk does increase as the consequence increases.

According to this graph, if the probability of an incident occurring is low and the consequence is low, then the overall risk to the organization will also be low. The company's risk response, therefore, may be minimal.

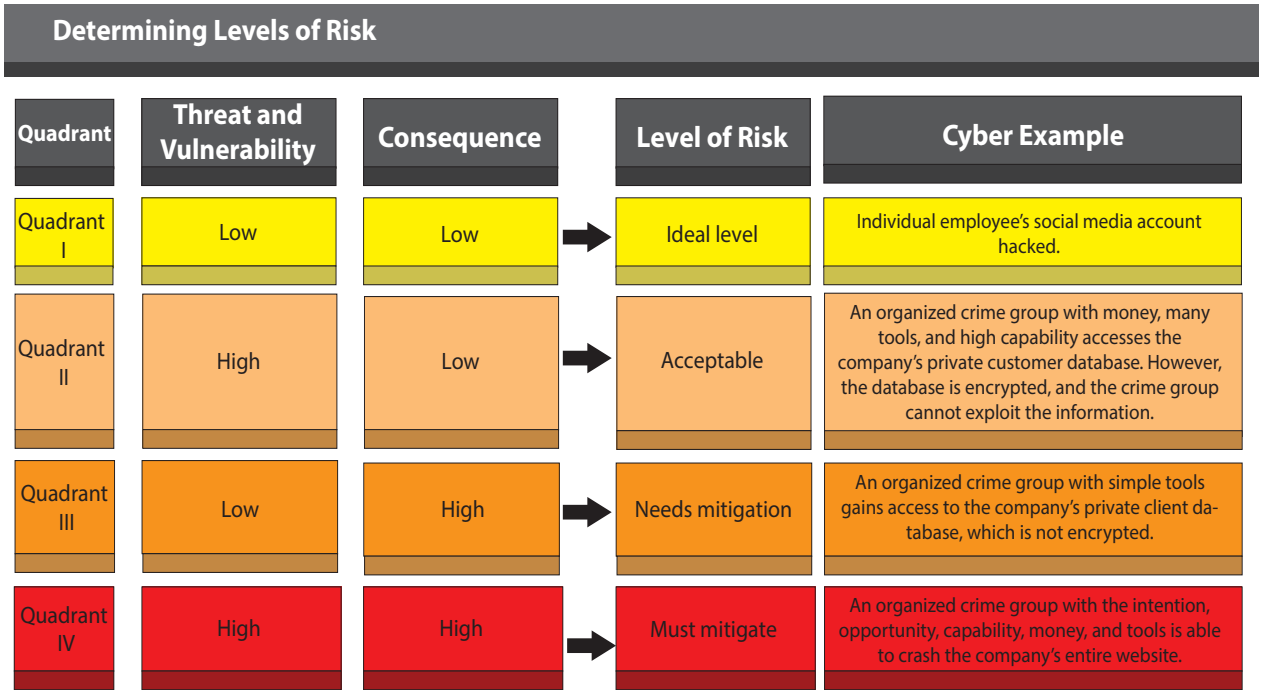

Figure 1-12: Risk determination quadrants.

However, if the probability of risk occurring is high, and the consequence of risk is also high, then the total level of risk is high and a comprehensive response plan must be implemented immediately. Figure 1-15 complements Figure 1-13, which explored the relationship between threat and consequence. Magnitude and severity is represented by Figure 1-13. The likelihood of risk is determined using a probability model represented by Figure 1-15. The graphs do not mirror each other, but when viewed and analyzed in conjunction, they reveal an in-depth and near-complete picture of risk.

Qualitative and Quantitative Risk Determination Methodologies

Determining the magnitude and the likelihood of risk is a process that can be conducted qualitatively, quantitatively, or using a combination of qualitative and quantitative methods.

Cyber Fact

Iran, China, the U.S., and Israel all have capabilities to launch offensive cyber attacks against other nations or groups. Both Iran and China are widely suspected of utilizing cyberwarfare tactics, while the Stuxnet worm organized and deployed by the U.S. and Israel in 2010 successfully neutralized Iranian centrifuges at an Iranian nuclear plant. Many countries have created national strategies for the prevention of cyber attacks, including the UK, with its 2011 Cyber Security Strategy. As part of the European Network and Information Agency, Germany, France, and the Netherlands have all implemented civilian cyber defense strategies.

Cyber Fact

Likelihood is the "chance of something happening, whether defined, measured or estimated objectively or subjectively, or in terms of general descriptors (such as rare, unlikely, likely, almost certain), frequencies, or probabilities," according to DHS's Risk Lexicon.

A **qualitative risk determination** approaches risk using non-numerical data. In a qualitative assessment, past records and data, patterns of data, behavior of technology, behavior of personnel, and interviews with personnel may all be sources that play into the determination of a risk's level and/or likelihood. A **quantitative risk determination** assigns numerical values to qualitative statements. In some cases, it is necessary to combine both qualitative and quantitative methodologies to best determine the level and likelihood of risk occurring.

All risk determination processes, whether surveys or statistical models, must begin with specific definitions of threat, vulnerability, and consequence that can form a risk research question. A **research question** or set of questions will guide the risk determination process. All research questions should be clear and concise and should accurately test the research subject. Though often extremely hard to do well, creating a research question or set of questions will provide a framework for the entire risk determination process; therefore, crafting research questions requires a fair amount of effort and thought.

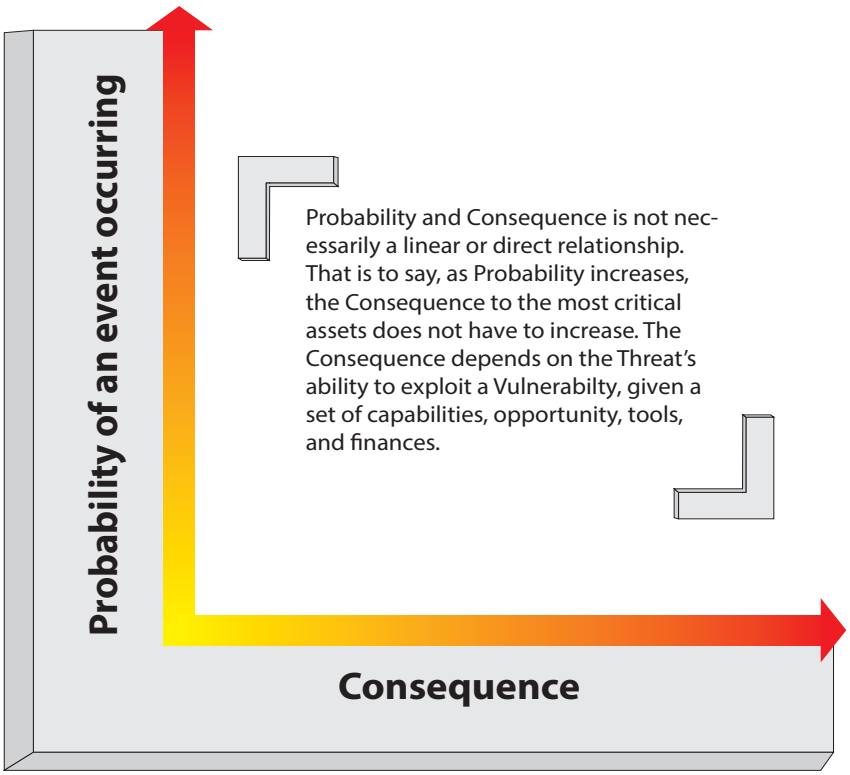

Figure 1-15: The likelihood of risk.

All cyber risk determination research questions should be valid and reliable to help standardize the process and to report the results of the risk determination. A valid design means you are measuring and/or testing exactly what you are interested in studying. Reliability indicates that if others were to conduct the same test through the same process, they would reach the same conclusion. In other words, **validity** indicates that the research answers the initial question, while **reliability** indicates the consistency of the research design. We will further discuss elements of research methodology in Chapter 6.

Qualitative and quantitative risk determination processes should follow the scientific method. The first step in the scientific method is to propose a hypothesis. The purpose of a hypothesis is to make a prediction or establish a relationship between two or more variables. It does not prove causality between one event and another; it simply states that a relationship exists between two variables. Hypotheses typically follow the model of an "if-then" statement, either explicitly or implicitly. Many researchers will begin by determining a **null hypothesis**, denoted H_0, which describes no relationship between variables. An **alternative hypothesis**, H_A, represents the relationship between the variables being tested. The goal of the research is to reject the null hypothesis in favor of the alternative hypothesis and prove a relationship between variables.

An example of a hypothesis is, "If one drives on New Year's Eve then there is an expectation of a car accident." Without an "if-then" statement, the hypothesis reads, "There is an expectation of a car accident when driving on New Year's Eve." Here is a null-hypothesis paired with an alternative hypothesis:

$\mathbf{H_0}$: An individual's IT security training has no influence in deterring an email spear phishing attack.

$\mathbf{H_A}$: An individual's IT security training is significant in deterring an email spear phishing attack.

Qualitative Risk Determination

Before jumping into the case study established by the hypotheses above, we will review the basics of qualitative research, particularly survey design. We will use the process of survey design to illustrate fundamental elements of qualitative risk determination.

Basics of Qualitative Design

Qualitative research is easy to do but hard to do well. The purpose of qualitative research is to define or explain descriptive events, personal views and understandings, or other kinds of ideological trends. Unfortunately, the flexibility of qualitative risk assessments often limits its precision. Unlike quantitative assessments, qualitative assessments may not be able to identify an exact risk measurement. Rather, they provide a more nuanced and holistic understanding of a particular issue.

Acquiring relevant qualitative data depends on the researcher's skills, available resources, and the research question. Data in qualitative designs can take the form of visuals, verbal discussions, and recorded observations. Interviews, focus groups, and surveys are popular methods of qualitative research, and all of these methods can be applied to the risk determination process.

Survey Design

A **survey** consists of a series of questions created in order to gather information about a particular research question. Although surveys are not always the most useful tools for determining exact relationships, they can provide descriptive or explanatory information about a large population that can later inform hypotheses about relationships between variables. The danger inherent in interpreting survey results is *extrapolation*—inferring a conclusion beyond the data.

Therefore, while quantitative research is generally better for generating precise numerical values related to risk, surveys allow management to gain insight into the level of risk that might exist within an organization. Surveys allow managers to understand a risk from the top to the bottom of their value chains because surveys are effective tools for gathering information across a population, particularly about human behavior, opinions, and preferences.

Surveys may take the form of personal interviews, self-administered surveys, or phone surveys. Each type of survey has its own benefits, drawbacks, and specific uses. For instance, personal interviews generally have high response rates and generate large amounts of detail from the respondents, but they are also the least efficient kind of survey to conduct in terms of time and money. Self-administered surveys are generally cost effective, but they have lower response rates. Lastly, phone surveys are cost effective, but may not generate the same depth of detailed information as a personal interview.

For more fundamentals of survey design for risk determination, see Appendix A.

Qualitative Example of Risk Determination

The National Nuclear Security Agency (NNSA) is a federal government agency that works to maintain the security and safety of the United States' nuclear weapons. This government agency is a likely target for a cyber attack because of the highly sensitive information about nuclear weapons that it is charged with protecting.

Let's look at a scenario in which there have been rumors that a foreign country with sophisticated hacking abilities and vast resources is attempting a spear phishing attack against the NNSA. The spear phishing attack will send an email from a seemingly known and friendly email address and will invite the recipient to click on a link in order to RSVP for a conference on radiation protection. When the recipient clicks, the hackers get control of the recipient's computer, thereby gaining access to the NNSA network and highly sensitive NNSA information.

Based on these rumors, the NNSA wants to test for its level of risk, specifically focusing on human vulnerability to a spear phishing attack. To determine its level of risk, the NNSA has decided to use a qualitative risk determination. The NNSA believes that this qualitative design will allow the agency to understand the human behavior that will determine the level of risk. Since the potential attack may come in the form of a spear phishing attempt, human error is one of the largest vulnerabilities to test for and work into the risk determination. Of all of the qualitative risk determination possibilities available, the NNSA has decided to design a survey to determine the level of vulnerability to such attacks within the organization.

The research question for the survey is:

"What is the level of risk within the NNSA to spear phishing email attacks targeting NNSA personnel?"

Cyber Fact

Phishing attacks are a popular form of cyber attack. These attacks "fish" for a user's personal information by trying to deceive the user into voluntarily entering personal information, or taking action that would provide the attacker with access to their system (e.g., "click this link" or "open this site"). **Spear phishing** is a targeted form of phishing—in order to carry out a spear phishing attack, the attacker must know something specific about the person being attacked. The spear phisher finds ways around common cybersecurity precautions by learning important facts about the person who is the target of their attack. They may learn what projects the person is involved with, or the people they communicate with most often. Spear phishing attacks are difficult to recognize because they are customized for the target and often look like they are coming from sources that the target knows and trusts.

The information security officer for the NNSA has decided that, for the sake of efficiency, she will conduct a self-assessment survey across a randomized sample of 150 NNSA employees. The NNSA employs approximately 3,000 people, so the sample will be roughly 5% of the total population. Using this model, the team of statisticians place each person's identification number into a random digit generator, thereby producing a sample group of 150 subjects to take the survey.

The information security officer has decided to distribute the survey via email. She chose this distribution method for two reasons. First, by having the participants fill out the survey on their computers, the data from their results will automatically be stored in a statistics software package. Second, in order to combat the low response rates that can accompany a computer-based survey, the security information officer requires that selected members print out a certificate designating that they have completed the survey. They are instructed to turn the certificate in to the Human Resources Department. Fortunately, this method is successful and results in a 100% response rate. Figure 1-16 displays the example survey. The purpose of this exercise is to illustrate the types of questions that the information security team would ask the personnel at the NNSA in order to evaluate the level of risk of a phishing attack. Before reviewing the survey, note the following six assumptions:

- All federal government employees have annual IT training.
- All federal IT training includes a module on phishing attacks.
- The NNSA conducts regular patching and vulnerability scans.
- Spam filtering is in place.
- .exe attachments are blocked in e-mail.
- Virus scanning is installed and kept up to date.

Based on these assumptions, the NNSA also assumes that cyber attacks against the organization will be highly sophisticated.

Survey for NNSA Personnel

Directions: Please fill out all questions to the best of your knowledge, using the following scale:

1= Poor

2= Fair

3 = Good

4= Excellent

5= Not applicable

Following completion of the survey, print out the certificate and return it to the HR Department.

Questions:

Training:

1. How informative was your most recent IT training?

a. 1

b. 2

c. 3

d. 4

e. 5

2. How well do you recall the information for your annual training?

a. 1

b. 2

c. 3

d. 4

e. 5

3. How well do you remember the phishing module from the annual training?

a. 1

b. 2

c. 3

d. 4

e. 5

Delivery:

4. How often do you receive emails from unknown sources?

a. Daily

b. Weekly

c. Monthly

d. Every few months

e. Never

5. Do you knowingly open emails that are delivered to your spam folder?

a. Yes

b. No

6. Do you pay attention to the addressee list of the email to verify the validity of content and/or sender of the email? In other words, do you pay attention to whether the addressee list makes sense given the nature of the email?

a. Yes

b. No

Payload:

7. How often do you open email attachments from unknown sources without verifying the email is valid?

a. Never

b. Sometimes

c. Often

d. Very Often

e. Not applicable/not sure

8. How suspicious are you of attachments in general?

a. Not at all

b. Somewhat

c. Suspicious

d. Very suspicious

e. Not applicable/not sure

9. If an error or malfunction occurs when you open an attachment (e.g., Adobe Acrobat crashes), do you ignore the error or check with the sender to receive a working version?

a. Ignore the malfunction

b. Check with the sender, and not inform the IT Department of the event

c. Check with the sender, and inform the IT Department of the event

Execution:

10. In general, do you report suspicious activity on your system to the IT Department?

a. Never

b. Sometimes

c. Often

d. Very Often

e. Not applicable/not sure

11. If your computer seems to be working harder than normal, do you report the abnormal behavior to the IT Department?

a. Never

b. Sometimes

c. Often

d. Very Often

12. If someone asks if you sent an email with an attachment but you did not, how likely are you to report this case to the IT Department?

a. Never

b. Sometimes

c. Often

d. Very Often

e. Not applicable/not sure

Other:

13. If you would like to elaborate on any of the questions or responses covered in this survey, please comment in the space below:

Figure 1-16: An example survey.

The NNSA will be able to analyze the results of this survey to determine the level of risk from a spear phishing attack that the agency faces. The NNSA has decided to aggregate the responses and generate the percentage for each response. The percentage for each response will inform the level of risk determined and therefore the risk response plan.

Quantitative Risk Determination

Now let's consider the quantitative process in risk determination. Quantitative analysis uses numerical values to test one variable, such as harm, against another variable, such as a threat's capability or lethality. There are many types of quantitative analysis methods that governments and organizations use in high-level risk determination processes. The following section assumes that the reader has a basic knowledge of statistics and logical reasoning. Statistics is a crucial element of quantitative risk determination and is therefore a foundation of the entire risk management process.

Probabilistic Risk Assessment (PRA)

The U.S. Nuclear Regulatory Commission (NRC), the Environmental Protection Agency (EPA) , and the National Aeronautics and Space Administration (NASA) all use **probabilistic risk assessment (PRA)**[7] in their risk management strategies. PRA is specifically concerned with:

- the magnitude or severity of the consequence, and
- the probability of an event occurring.

A PRA model generates a value or range of values for the probability of risk occurring. The probability is based on a set of variables. **Variables**, also known as *indicators*, are the various factors that influence the likelihood of risk occurring. For example, in the New Year's Eve scenario, the variables can include: the safety features of the car, the time of night at which the friends are on the road, the number of people in the car, and the alertness level of the driver. All of these factors influence the likelihood of getting in an accident or not getting in an accident on New Year's Eve. Indicators of a burglar invading a family's home include whether or not the front door was locked, whether or not a fence surrounds the perimeter of the home, whether or not an alert system is installed, and whether or not there are any witnesses present when the burglar attempts the invasion.

In a PRA model, consequences are expressed numerically. In the New Year's Eve example, we will need to consider the number of people potentially injured. In the cyber example, we will need to consider the number of hours that the website is defaced and the number of customers affected. Our goal is to graphically represent the probability of an event occurring; this probability is the unknown value. Event trees, multivariable statistical analysis, and other PRA models generate equations or processes to determine a value for the probability of a risk occurring. We will return to specific PRA models and a statistical model later in this chapter. First, we need to have an understanding of the quantitative process and probability.

Basics of Probability

Probability can be as simple as flipping a coin or as complex as calculating the probability of catastrophic damage to New York City's electrical grid. When studying probability, the following questions arise: Will an event occur? How certain can I be that an event will occur?

Probability (P) is defined as a value between 0 and 1, representing 0% and 100% probabilities that an event will occur. You cannot have a 150% chance of an event occurring, just as you cannot have a -50% chance of an event occurring. Therefore, probability distributions are between 0 and 1. The sum of all of the probabilities for a given scenario must equal 1, or 100%. Figure 1-17 illustrates probability distribution.

$0 \leq P(\text{event}) \leq 1$

Figure 1-17: Probability distribution.

Flipping a coin has two possibilities, heads or tails, so you have a one out of two probability (50%) of landing heads. You also have a 50% probability of landing tails. In a deck of cards, there are a total of 52 cards, of which exactly four are queens, so the probability of randomly picking a queen from a deck of cards is 4/52 (roughly 8%).

Type of Probability	Example
Independent	Flipping a coin
Mutually Exclusive	Rolling a dice
Not Mutually Exclusive	Drawing a card out of a normal deck of cards
Conditional	Event A depends on Event B (expressed as "the probability of A given B")

Figure 1-18: Definitions and examples of probability.

The terms **independent** and **exclusive** do not have the same meaning in probability as they might in everyday language. The term *independent* suggests that the knowledge of one event occurring does not affect the probability of the other, or next, event occurring. Rolling a seven on one throw of the dice, for example, does not diminish or increase the probability that the next throw will also be a seven. *Exclusive*, also referred to as disjoint, indicates that two events can never occur at the same time. **Conditional probability** is defined as the probability of one event occurring with the knowledge of the outcome(s) of the other event(s).

Our study of cyber risk management is going to focus on conditional probability, because the likelihood of harm is contingent upon multiple events either occurring or not occurring. Once we understand the basics of probability, we can use this knowledge to perform research using PRA models or statistical models.

As discussed earlier, we need to set up a null and alternative hypothesis to help us solve for risk. A null hypothesis, denoted H_0, describes no relationship between variables. An alternative hypothesis, H_A, represents the relationship between the variables in research based on the scientific method. The goal is to reject the null hypothesis. In an "if-then" statement, the H_A asserts: "If malicious code is sent to a company, then the network will be harmed." To be clear, this statement does not indicate causality between events. For example, malicious code alone does not cause information to be leaked.

We will test the following hypotheses using different quantitative models.

$\boldsymbol{H_0}$: There is no association between malicious code and harm to a company's website.

$\boldsymbol{H_A}$: There is an association between malicious code and harm to a company's website.

The Four PRA Models

Four different PRA models can reveal a level and likelihood of risk. The following section provides examples for event tree analysis and fault tree analysis.

Event Tree Analysis

Event tree analysis traces the probability of a response to an incident. This response is usually expressed as a binary—yes or no.

Event trees create great visuals for tracing a sequence of events and the probability of each event. Event tree analyses begin with an initiating event. In our New Year's Eve example, the initiating event is the drive to the nightclub

or party. In the cyber example, the initiating event is malicious code being sent to an online company. An event tree then traces the possible outcomes after the initiating event occurs and indicates the probability of subsequent events.

In an event tree analysis, an outcome is the probability of a combination of events occurring. If we were to make a chart for the possibility that a car speeds in the New Year's Eve example, we might also take into account some other potential outcomes, such as skidding on ice and brakes failing. We would then continue to trace outcomes back to the initiating event. "What will happen if the car speeds, skids on ice, and brakes fail?" is one sequence of events that will produce a unique outcome. After the event tree is complete, each sequence will be quantified and the outcome of each event happening will be calculated, thereby quantifying the potential risk of any given incident on the event tree.

Figures 1-19 and 1-20 show event trees for two scenarios, the New Year's Eve drivers and the malicious code sent to an online company. The "Risk Level" boxes at the rightmost end of each tree refer to the four risk quadrants that were depicted in Figure 1-12.

In each case, the risk level is based on the severity of the sequence of events. In the New Year's Eve scenario, for instance, speeding, skidding, and brake failure constitute a more serious and dangerous sequence of events than not speeding, not skidding, and normal brake function; hence, the first sequence results in a Quadrant 3 (Q3) assessment, while the second generates a Q2. Similarly, in our cyber example, if an intrusion detection system is not in place, if a firewall is not protecting a server, and if there is no policy to discourage employees from opening suspicious emails, then the risk level is greater than it would be in a scenario in which all of these elements had been implemented. Based on the definition of conditional probability, the final outcome in each case is calculated by multiplying the probabilities of each incident.

Fault Tree Analysis

Fault tree analysis examines an entire organization or enterprise from a top-down perspective and identifies the combination of failures that may contribute to the system's malfunctioning; it is a deductive methodology, because it derives possible conclusions from a given set of premises.

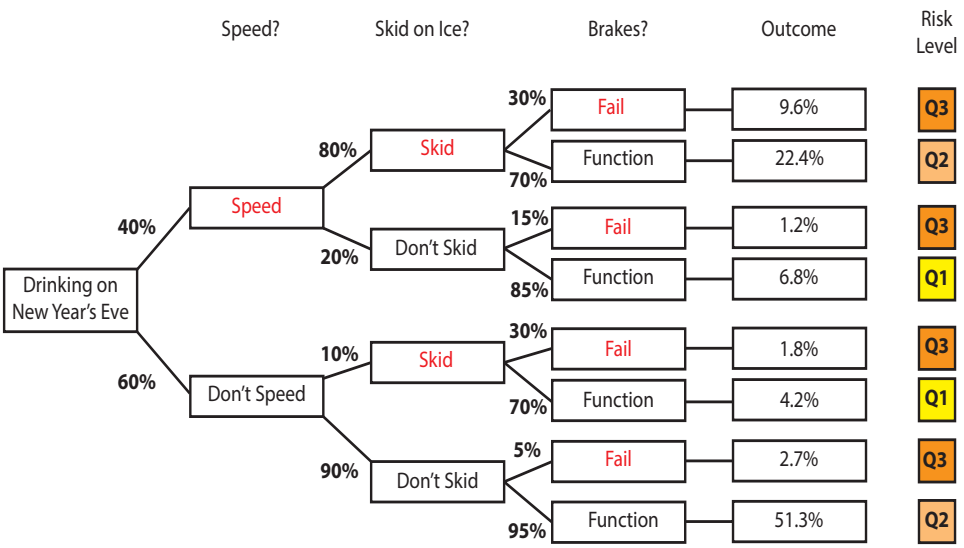

Figure 1-19: Event tree for the New Year's Eve scenario.

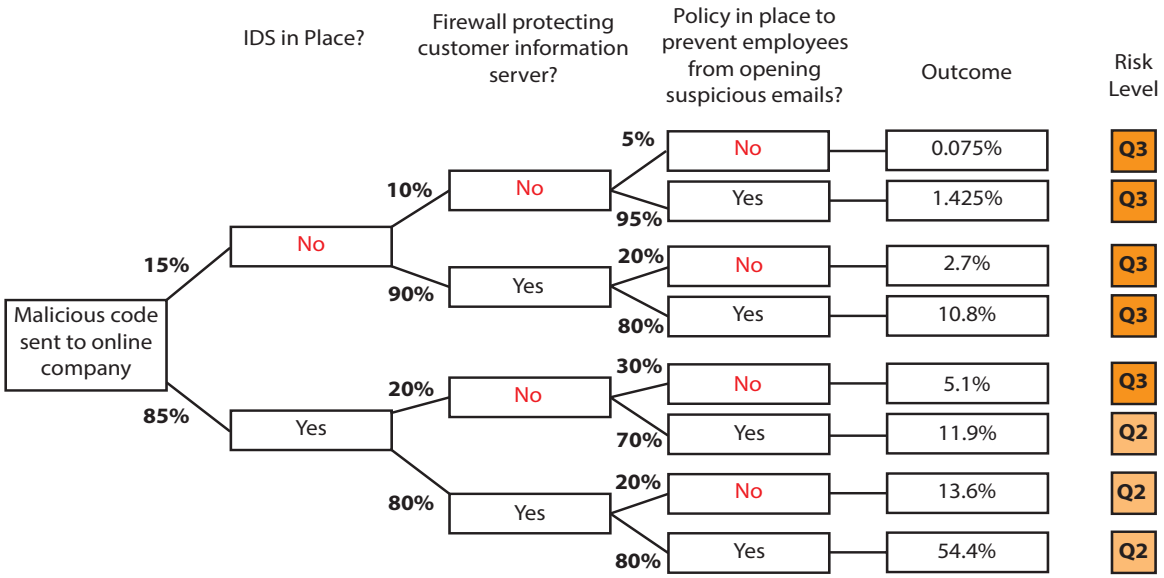

Figure 1-20: Event tree for the cyber scenario.

Risk analysts and risk managers use fault tree analysis to analyze the probability of an undesired final event. Using this type of analysis, an upper level manager begins by designing the top of a fault tree diagram—beginning with the undesired event. This is where a fault tree analysis would place an event such as a car crash in the New Year's Eve scenario, a burglar intruding into a family's home, or a malicious code taking down a website. In fault tree analysis, the top events are always the final events.

From a final event, the risk analyst or risk manager will trace possible causes of that event using a series of coordinated symbols. These possible causes could be traced from one particular event that begins the line of action until reaching the final event at the top of the fault tree diagram. The symbols and their meanings are shown in Figure 1-21, while Figures 1-22 and 1-23 show fault tree diagrams for the New Year's Eve and cyber scenarios.

Human Reliability Analysis (HRA)

Human reliability analysis (HRA) calculates the human errors that may affect a risk-induced incident. The probability of a risk occurring is determined using indicators that can influence the likelihood of the risk occurring. Examples of indicators are the sophistication of training for personnel and the implementation of policy procedures for an individual or an organization. In the New Year's Eve case, the possibility for human error is particularly important because, more than on any other night of the year, a large number of intoxicated drivers are likely to be on the road. The more intoxicated a driver is, the more likely it is that a human error will occur. Therefore, to calculate the risk of driving to and from the nightclub or party, it would be helpful to use an HRA analysis to quantify the amount of human error that factors into the risk of driving on New Year's Eve.

In the cyber realm, spear phishing techniques explicitly aim to exploit an individual. In this way, spear phishing capitalizes on human vulnerabilities.

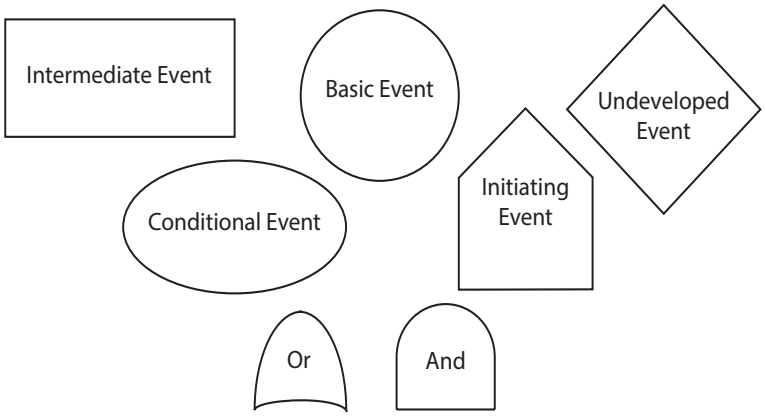

Figure 1-21: Fault tree symbols.

Monte Carlo Analysis

A **Monte Carlo model** is a multivariate statistical analysis for computing mathematical risk. A Monte Carlo analysis considers possible variations in each individual factor of the analysis and examines the many possible ways in which the factors may interact. For example, one set of factors may influence the likelihood of risk occurring to a greater or smaller degree than another set of factors. A Monte Carlo method will test thousands of iterations of ways that situational factors may combine to create a simulation of a scenario with a large enough sample size of possible outcomes to calculate probability statistics for each outcome.

For the cyber risk determination, a Monte Carlo model would take into account factors such as the implementation of a firewall, updates to anti-virus software, encrypted hard drives, properly secured computers, and threat

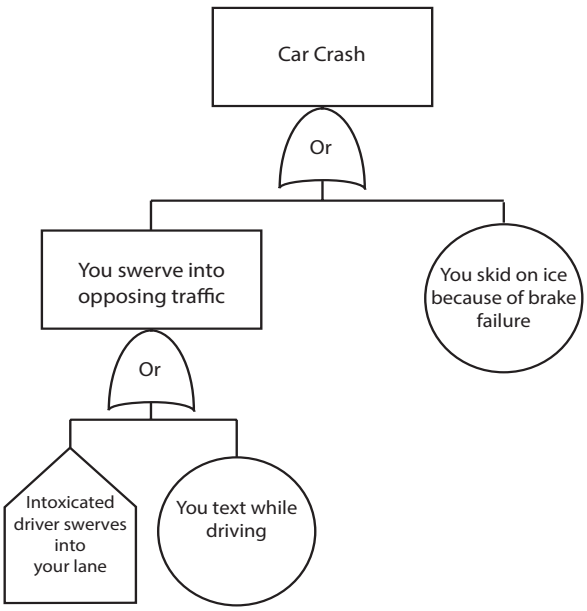

Figure 1-22: Fault tree for the New Year's Eve scenario.

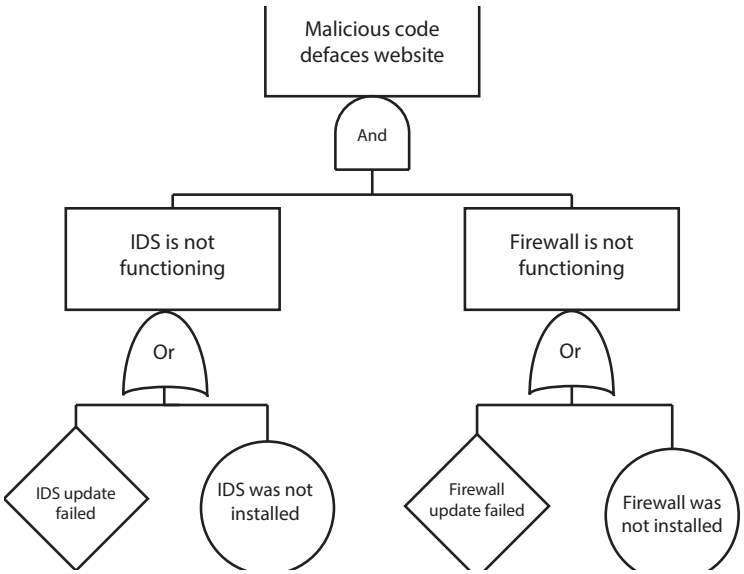

Figure 1-23: Fault tree for the cyber scenario.

levels to determine the risk to an organization's computer network. The model would then run simulations varying those factors and calculate the overall risk, as well as the variations to the risk that each factor presents.

The Monte Carlo model generates a probability value that a risk will occur. This model will inform the organization about the factors that contribute most to the risk of harm to their network, so that they can take steps to reduce these contributing risk factors as much as possible. The Monte Carlo method is not the only statistical model used to identify the probability of an event occurring, but it is one of the most popular models that professional risk managers use.

These PRA models represent only a few of the quantitative methods available to risk managers to conduct in-depth risk determination processes. Risk managers, like many other social scientists, have started to use a wide variety of mathematical models to quantify their reasoning and to provide evidence for their hypotheses. These models include game theory, logic decision processes, and multivariable statistical analysis, which we will discuss later in this chapter. The appropriate model to use in any given risk management situation depends on the kind of answer a risk manager is seeking and the type and amount of data available.

Statistical Modeling

An online retailer wants to reduce the risk of customer information being leaked; the retailer, therefore, is interested in which security controls are the best at preventing information from being leaked. The mathematical concepts and statistical background helpful to understanding risk are available in Appendix A.

A Quantitative Risk Determination Example: Determining the Likelihood of Harm Occurring to an Online Retailer

The following probability model is for illustrative and educational purposes only. The dataset was randomized using Microsoft Excel in order to illustrate how risk can be quantified using statistical inference. All statistical results were produced using Stata 11.

Let's imagine that the risk management team at a popular online retailer wants to know what kind of risk they face with regard to critical customer information being leaked. They want to answer the following question:

"Which security controls are the best at protecting customer information from being leaked?"

To evaluate this question, we will use a statistical regression model. A **regression model** uses multiple independent (*x*) variables to predict or estimate a dependent (*y*) variable. Presumably, a change to any of the independent variables will affect the outcome(s) of the model. In other words, the *y* value is dependent on the *x* value(s). In the case of multivariable analysis, each *x* value must be mutually exclusive; the *x* values cannot be affected by, nor dependent upon, each other.

Cyber Fact

Game theory is a strategic decision making process between rational actors. The "prisoner's dilemma" is an example of a game theory model. DHS's Risk Lexicon defines game theory as a "branch of applied mathematics that models interactions among agents where an agent's choice and subsequent success depend on the choices of other agents that are simultaneously acting to maximize their own results or minimize their losses."

In order to determine which security controls are the best at securing information, 200 (n=200) online retailers will be examined. They will answer the question of whether or not, yes or no, customer information was leaked on one random day (24 hours) in 2011.

Because we have a binary (yes or no) question, we will use a **logistic (logit) regression model** to evaluate risk. Logit models are used when the dependent variable is a binary response. The model generates a probability value. A logit model reveals the most statistically significant indicators and the type of association the indicators form with the dependent variable, the latter expressed as *positive* (direct) or *negative* (indirect). A positive relationship indicates that if the independent variable increases, the dependent variable increases; a negative relationship means that if the independent variable increases, the dependent variable decreases (or vice versa).

Analysis of the Most Significant Factors Influencing Customer Information Leaked

Based on the most complete model generated using the logit equation (Model 3), *Strength of Password* and a *2-Factor Authentication* are the two most significant security controls.

An example of a strong password would be one that includes capitalization, a number, a symbol, and at least eight characters unrelated to any personal information; for example, "rh7TL!9" would be a strong password. A particularly weak password might be a word that is found in a dictionary, such as the word "password."

A 2-factor authentication is a security control that requires the user to enter two pieces of data from the following three: "something the user possesses" (such as a member number); "something the user is" (such as an item of biometric data); and "something the user knows" (such as a password).

According to the logit analyses (see Appendix A for the logit analyses of Models 1, 2, and 3), the independent variable *Password* is significant in determining the probability of a risk, defined as customer information being leaked, and as the password becomes stronger by gaining in complexity, the likelihood decreases that customer critical information will be leaked.

Analysis, however, shows that a 2-factor authentication is the most significant variable, based on its low p-values in all three models.[8] Therefore, if a 2-factor authentication is present as a security control on a company's website, the likelihood of customer information being leaked decreases. According to this analysis, and based on its statistical significance in each model, 2-factor authentication is the most significant form of security control that the online retailer can use.

Interestingly, in Model 3, *Number of Users* was another statistically significant variable.[9] This result indicates that as the number of users increases, the likelihood of customer information being leaked decreases. Risk analysts may draw many different conclusions from this outcome. When we consider this statistically significant variable, it is important to remember that correlation does not equal causation; the number of users does not cause a website to be defaced or not be defaced. Rather, and most likely, the most popular online stores, measured by number of users, must have the greatest security controls in place for their consumers. Therefore, these stores also have a lower probability of customer information being leaked.

The *Strength of Hacker* variable was only significant in Model 3. Therefore, common knowledge prevails: as the strength of the hacker increases, the likelihood of customer information being released increases.

To sum up the findings from this study: When an online retailer considers which security controls to employ in order to protect customer information from being leaked, requiring a strong password of its consumers and consideration of a 2-factor authentication should be at the top of the list.

Based on these models, risk managers may suggest to the chief information officer (CIO) or chief information security officer (CISO) of an online retailer that the website require a strong password and a 2-factor authentication in order to secure customer information. Decisions made in these areas should be based on an effective risk management program. However, the CIO and/or CISO have limited resources for the deployment of security controls, so they must choose a combination of controls that securely safeguards customer information, avoids customer inconvenience, and works for the company's budget. At this point, a cost-benefit analysis is necessary to determine which controls should be in place.

Predicting the Probability of Customer Information Leaked

The strength of multivariate statistical analysis lies in its predictive capabilities. A thorough, valid, and reliable design will allow a risk manager to extrapolate future trends that will help the company to protect its critical assets. A CIO or CISO's job is to present the company's board of directors with the best set of security controls, given limited resources and other constraints, such as the constraint of maintaining a user-friendly website. The specific mathematical forecasting techniques used to predict the probabilities of customer information leaked based on this scenario are available in Appendix A.

The graph in Figure 1-24 expresses the probability of customer information leaked as the "Probability of Risk Occurring." One can see how the probability of customer information leaked changes as the degree of security controls differs in each model.

Cyber Fact

The **CIO** is in charge of an organization's IT planning, budgeting, and performance, including its information security components.

Recall that the best security controls for protecting customer information are a strong password and a 2-factor authentication. However, the risk analysis team wants to know the combination of security controls that will best protect customer information from malicious code, based on the resources allocated and constraints placed on the online retailer.

Model A is the riskiest scenario: There is a 93.136% likelihood of customer information being leaked under the set of conditions in Figure 1-24. Model D is the safest and most realistic option, which even takes into consideration a hacker with strength level 7 out of 10, 10 being the strongest.

Model D is the only model that has a 2-factor authentication, and it has the strongest password. While Model D does contain multiple security controls, it is a realistic option. Consumers may not enjoy having a 2-factor authentication; however, on a popular website that is a target for strong and capable hackers, consumers should feel confident that there is only a 0.641% likelihood that their critical information may be leaked. There are additional models to reduce the probability of risk occurring, but incorporating additional levels of security may become too burdensome for the consumer and deter them from using the site. The store must strike a balance between a high level of security and consumer convenience.

The likelihood of a risk occurring and the level of risk are two of the most crucial pieces of information that go into the risk management process. Determining the likelihood of a harm (in our example, customer information leaked) occurring is a critical calculation for risk managers. Following our risk determination process in this example, the CIO or CISO now has four different models to present to the board of directors as reasons to promote the security controls of Model D versus Model A. They will emphasize the need for 2-factor authentication and a strong password for consumer login. At the same time, a good risk analyst, CIO, or CISO will account for the possible variations in a potential hacker's level of sophistication, as well as other external factors that play a role in the outcome of the probability model. Variation in the independent factors will prepare the company for a variety of situations. For these reasons, the risk determination process is one of the most important phases in the formation of an effective risk response. By quantifying and identifying various and varying risks, a manager will be well prepared to respond to a dynamic risk environment.

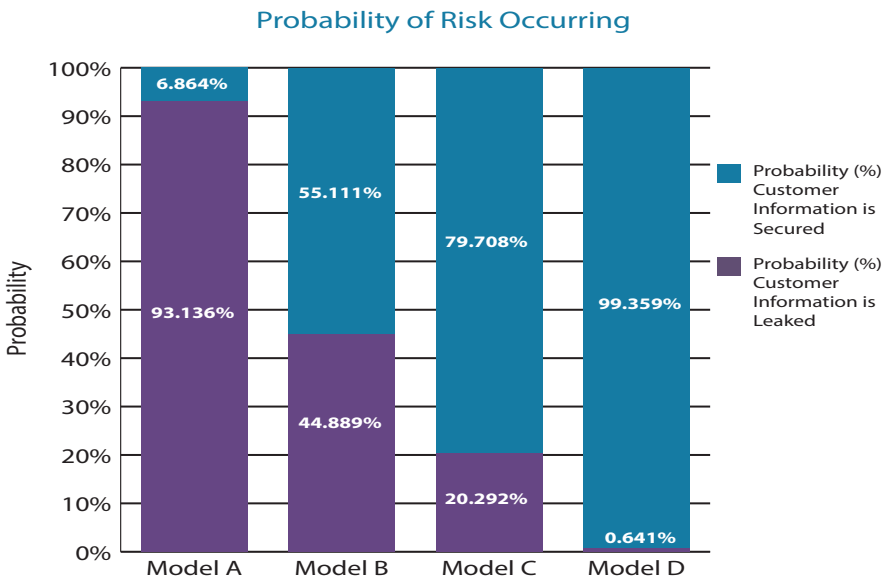

Figure 1-24: Graph of Probability of Risk Occurring.

Risk Response

A **risk response** is the most appropriate response to the determined risk. After a thorough evaluation of many possible responses, the response that the risk managers choose prevents harm to an individual, organization, or enterprise, while also staying within budget. Potential responses depend upon the threat's capability and the criticality of the assets. Risk managers evaluate these potential courses of action against the calculated risk, and then choose the most appropriate course of action. A risk response plan may incorporate various mechanisms, decisions, or policies to construct the best response to protect against risk.

The response plan must fit the determined level of risk; the course of action for a high-risk situation is not appropriate for a low-risk situation. For example, installing a firewall at the FBI against the threat of a foreign adversary hacking into the computer system is not the appropriate course of action to protect this high-risk and nationally critical asset.

In a low-risk cybersecurity situation, a risk response plan may be as simple as encrypting a hard drive. However, if the cyber risk is determined to be high, then the risk response plan may be more complex and include buying additional monitoring technicians, purchasing cyber insurance, replacing virus software, and creating stronger passwords for employee computers. Regardless, the risk response plan will be constrained by the available resources.

In short, the purpose of the risk response is to evaluate, determine, and implement the best course of action to contain and manage the determined risk. A risk response plan *accepts*, *avoids*, *mitigates*, or *transfers* risk. Figure 1-25 outlines these four different types of risk response plans and applies them to our cyber example.

Risk acceptance does not mean that one allows a threat to cause damage. Rather, it means that one accepts that a certain degree of risk is unavoidable. An example of accepting risk is allowing one's email system to permit any user to send an email. When one accepts this scenario, one accepts the ever-present risk that an email may contain malicious code that will infect a computer—and perhaps a computer network; however, email correspondence is essential for conducting business, so it makes sense for a company to accept this risk. To refer back to our New Year's Eve example, the group of friends accepts risk by riding in a car at all, and accepts even more risk by riding in a car on New Year's Eve.

Risk avoidance is perhaps the most obvious way of reducing the possibility of harm. Risk avoidance means that one evades or circumvents a threat by removing a critical asset from potential harm. However, avoiding risk does not mean eliminating risk, and risk avoidance is rarely a practical strategy for conducting business or living

3.

Risk Response	
Accept	Receive emails from unknown users.
Avoid	Stop using the company's computers.
Mitigate	Protect the server with a firewall.
Transfer	Purchase cyber insurance.

Figure 1-25: Risk response elements and examples for an online retailer.

one's life in a risk-filled world. In the New Year's Eve scenario, avoiding risk would mean staying at home and not traveling anywhere. In our cyber example, avoiding risk could mean not using computers at all to conduct company business. Of course, both of these risk-avoidance measures are highly inconvenient to people who want to celebrate a holiday or conduct business in the modern world.

Risk mitigation is a strategy to contain an imminent or current incident. In the cyber example, risk mitigation includes taking steps to manage the protection of network operations and recover damages when vulnerabilities are exposed or cyber incidents occur. An example of a company's mitigation device might be placing a firewall on a network when a threat is identified. The New Year's Eve partiers may mitigate the possibility of a potential crash by assigning a designated driver. The family can mitigate risk by installing an alarm system.

Risk transfer is a strategy that rests on the assumption that although risk can never be eliminated, it can be displaced or shifted. The DHS Risk Lexicon defines risk transfer as an "action taken to manage risk that shifts some or all of the risk to another entity, asset, system, network, or geographic area." An example of risk transfer is the purchase of insurance.

Cybersecurity insurance works the same way that automobile or home insurance does: it transfers some of the risk of operating in cyberspace. For example, cybersecurity insurance protects a company against some of the fallout in case the company suffers a data breach. Insurance may cover the fees associated with notifying customers of a breach, including consulting costs and credit and fraud monitoring services. Other cybersecurity insurance policies may also cover legal expenses. Insurance does not eliminate the risk of doing business in cyberspace, nor does it cover all of the direct and indirect consequences that may stem from a cyber emergency, but it does make cyberspace a more friendly and attractive place to do business, in spite of the risks.

Risk Monitoring

The final stage in the cyber risk management process is **risk monitoring**. Risk monitoring is an ongoing process that takes place after the response plan has been implemented. The risk monitoring stage involves developing a strategy for constantly monitoring changes in threats, vulnerabilities, and criticalities. There are three main elements in the risk monitoring phase: compliance, effectiveness, and identifying changes (see Figure 1-26).

Compliance measures how well an implemented policy has been followed. If a policy has been issued, then it is the duty of the monitoring team to guarantee that the policy is followed. For example, in the cyber-related scenario, compliance means verifying that the prescribed firewall has actually been installed. The next step is to guarantee that the firewall is working—in other words, to check and ensure its effectiveness. If it is not working, then it must be replaced, or an alternative protection mechanism must take its place. The risk monitoring process also takes note of changes to the system or the risk environment. This process of identifying changes may alter the way the risk is framed and therefore start the risk management process all over again. This step in the risk monitoring process incorporates all of the information within the risk determination and response sections.

A company may identify changes based on an intelligence report. For example, the company may learn from the U.S. government that a foreign adversary has invested additional money in its cyber warfare capabilities. This foreign adversary now represents a threat to the company. Since the threat has grown in strength, the company will have to reevaluate its risk management strategy to account for the new level of threat.

Figure 1-27 provides a breakdown of each aspect of risk monitoring with an application for our cyber example.

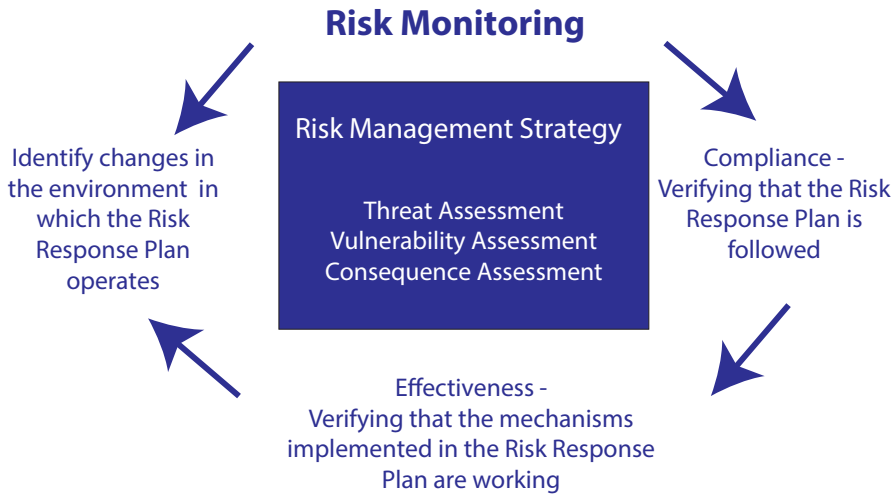

Figure 1-26: Risk monitoring cycle.[10]

Monitoring risk is an ongoing process of reevaluating and refocusing upon the threats, vulnerabilities, and potential harm to critical assets. The monitoring process should methodically identify risks, manage the response plan, and collect data and information for future assessments. The risk monitoring process will also shed light on the best allocation of resources, and whether or not a redistribution or reprioritization of resources should take place.

Conclusion

Risk management is a fundamental process safeguarding the critical assets of both individuals and industries. The formalized structure of the risk management process provides the regimented focus required for generating the most comprehensive risk assessment and therefore the most appropriate risk response.

There are multiple methods for performing the risk determination. This chapter highlights only a few ways of determining risk: event and fault trees, qualitative surveys, and statistical multivariable analysis. The rigid scientific method of calculating risk can be applied to the risk determination process no matter what method is used to quantify the risk.

In some cases, taking on risk can be good for businesses and individuals. Taking on risk may allow individuals and/or organizations to advance their positions competitively, and taking on risk may be the only way to try something new. However, this chapter focuses on how the negative repercussions of an unintended risk can cause significant harm or damage.

Because of the cyber realm's many nuances, challenges, emerging developments, and inherent complexities, it is helpful to approach the protection of cyber assets through a risk management lens. By applying the risk management process identified in this chapter to cybersecurity, individuals, companies, governments, and other enterprises will have: (1) identified their most important assets; (2) evaluated the threats to the assets, the vulnerabilities that a threat could exploit, and the consequences if these vulnerabilities were succesfully exploited; (3) determined the level of risk and the likelihood of risk occurring, using either qualitative or quantitative reasoning; and (4) generated a risk response to counter the threat.

4. Monitor Risk	
Compliance	Verify the implementation of a company-wide firewall.
Effectiveness	Continually test and ensure that the firewall is working.
Identify Changes	Identify new threats and vulnerabilities, critical assets, and consequences as the risk environment changes and evolves.

Figure 1-27: Risk monitoring elements.

Risk management is a circular process; it never stops. All cybersecurity professionals must approach risk management as a key aspect of their job. Risk management is not merely a compliance issue to be considered once a year. Rather, every organization's cyber risk management plan must also be refocused and adjusted constantly to protect against the latest threats in cyberspace.

Key Questions:

1. Chart the steps of the risk management process for two scenarios:
   (a) The scenario of protecting the most critical asset in your home.
   (b) The scenario of protecting the most critical asset, sensitive documents, on a government computer.
2. How might the private sector and public sector approach risk management differently? How might their approaches be the same?
3. Which of the following is NOT an element of risk framing?
   (a) Tolerance
   (b) Trade-offs
   (c) Threat assessment
   (d) Constraints
4. Which of the following is NOT an element of risk response?
   (a) Avoid
   (b) Mitigate
   (c) Manage
   (d) Transfer
5. What does the PRA model focus on?
   (a) Magnitude and severity of risk
   (b) The consequence assessment
   (c) The probability of risk occurring
   (d) Risk response

6. Which is an effective method of determining the level and likelihood of risk?
   (a) Human reliability analysis
   (b) Statistical modeling
   (c) Event tree analysis
   (d) Survey design
   (e) None of the above
   (f) All of the above

Cyber Connections

Some individuals, companies, and organizations view risk management as a burden or a checklist merely to comply with standards and regulations. This approach fails to consider that a clear and robust risk management strategy is essential to any successful program, endeavor, business, or institution. See Chapter 3 for more on program management. See Chapter 5 for more on the legal obligations of corporations with regard to risk management.

A robust risk management strategy means a strategy in which all elements of the specific risks an organization faces, as well as the organization's most critical assets, are explicitly understood by all stakeholders. The strategy should not be a checklist of standards to be checked off in order to meet the demands of an outside regulator. Rather, the organization's risk management strategy must combat identified risks with a specific plan that will be continually monitored and reassessed.