CHAPTER

14 Risk Analysis, Incident Response, and Contingency Planning

RISK MANAGEMENT IS AN IMPORTANT information security tool. The risk management process helps an organization understand the risks, vulnerabilities, and threats that it faces each day. It helps the organization understand its security posture. It also helps the organization know where to strengthen that posture. An organization can't meet its information security goals if it doesn't understand its risks. It may not be able to properly protect its resources and data.

This chapter focuses on information technology risk management. It reviews fundamental risk concepts and how they're applied. It explains how organizations use risk management to help them create their other contingency plans.

Chapter 14 Topics

This chapter covers the following topics and concepts:

- How to plan for contingencies
- What risk management is
- What three types of contingency planning are
- What incident response is
- What some special considerations are

Chapter 14 Goals

When you complete this chapter, you will be able to:

- Describe the risk assessment process
- Describe how to create an incident response plan
- Describe the business continuity planning process
- Describe how to create a disaster recovery plan
- Explain the differences between risk assessment, incident response, business continuity planning, and disaster recovery

Contingency Planning

Organizations must plan for many events. Contingency plans don't focus just on an organization's information technology (IT) assets. Instead, the field has grown to include all types of planning to make sure that an organization can continue to operate in the event of an emergency or disaster. In recent years, the United States has seen several dramatic events that highlight the need for all types of contingency plans. The September 11 terrorist attacks of 2001. The northeastern U.S. power-grid failure in 2003. Hurricane Katrina in 2005. The H1N1 flu outbreak in 2009.

When organizations talk about contingency plans, they're talking about holistic plans. These are plans that cover the whole organization and all of its processes. They identify operations that are critical to the business's survival and recovery. These processes include IT resources and operations. Many organizations store their data and records electronically. Many of them have grown dependent on their IT resources and the automated processes those resources provide. Thus, focusing on how to protect and recover IT assets is an important contingency planning component. Some studies indicate that almost half of all companies that experience a data loss never recover.[1]

This chapter discusses contingency planning processes as they relate to IT resources. The different planning processes discussed in this chapter are:

- Incident response planning
- Disaster recovery planning
- Business continuity planning

The scope of any kind of contingency planning is very broad. This chapter focuses on only one narrow piece: planning for the security of IT resources. While all contingency plans have different goals, the foundation for these processes is the same. In order to prepare contingency plans, organizations must analyze and plan for the IT risks that they face. The risk analysis and foundation process helps an organization understand how it needs to protect its IT resources. An organization can't make any contingency plans until it understands the risks to its IT resources.

FYI

This chapter talks about risk management, incident response, and contingency planning. Risk management and contingency planning are used to protect IT resources. Incident response is a type of contingency planning that an organization uses to react to attacks against its IT infrastructure. Disaster recovery and business continuity planning are contingency plans that help an organization continue business operations following a disaster.

It's important to remember that protecting IT resources isn't the only goal of a contingency plan. It's not even the most important goal. Natural and man-made events disrupt thousands of lives each year. The most important goal of any type of contingency plan is to preserve human life. Continuing operations and restoring data are secondary goals.

Risk Management

The National Institute of Standards and Technology (NIST) says that risk management is "a complex, multifaceted activity that requires the involvement of the entire organization."[2] **Risk management (RM)** helps an organization identify the risks that it faces. It also makes sure that organizations respond to risk in a cost-effective manner. Organizations use RM to support their business goals.

One of the main goals of RM is to protect the bottom line. When risk is realized, it negatively affects an organization's profits. RM helps an organization align its information security needs to its business goals. It makes sure that an organization spends its limited resources wisely and in ways that enhance business goals. An organization uses RM to plan for information security.

For example, an organization has a database that holds customer data. The database is critical to the organization's business. It uses the information in the database to develop products for its customers. The database also holds marketing information. The organization would be harmed if it didn't have access to this resource. Its development and marketing activities would be negatively impacted. The organization must protect this database.

A risk analysis shows that the database is at risk because system administrators don't use strong passwords. This is a vulnerability. As a result, the database is open to attack. If it's attacked, the confidentiality and availability of data on the resource are compromised. Attackers could steal this data. They also could harm the database so that the organization couldn't use it. The risk analysis shows that the vulnerability (weak passwords), together with the threat (attackers), is very critical given the organization's business.

As part of its RM function, the organization must review the potential impact of this risk. The cost of a data breach can be quite high. The organization might have to notify its customers about the breach. It also might have to report the breach to its regulatory authorities. Those agencies could fine the organization for failing to secure its data. Its customers could sue it as well. The organization could lose customers. It also faces its own operational issues if that database isn't available for a certain period.

Once the organization identifies the risk and potential impact, it can take steps to mitigate it. In this example, the organization must reduce the risk posed by weak passwords. It can require its system administrators to use strong passwords. It also can configure its systems to enforce password complexity requirements technically. Finally, it can make sure that it backs up its customer data properly. In this short example, the RM process identified a risk to a resource. It also made sure that the organization addressed that risk. The company was able to secure an important resource. As a result, it was able to protect its business operations and bottom line.

Each step in the RM process supports an organization's business goals. The most basic RM process includes the following steps:

FYI

The RM process described in this chapter is very similar to the risk management framework used by the U.S. federal government under the Federal Information Security Management Act (FISMA). An organization could easily use that framework to protect its IT systems. The National Institute of Standards and Technology has created guidelines that federal agencies use to implement this framework.

- **Risk assessment**—Identify the threats and vulnerabilities to the organization's IT resources. Determine the impact of those threats and vulnerabilities.
- **Risk response**—Use policies and controls to respond to risk. An organization responds to risk according to its business strategy.
- **Employee training**—Train employees on known threats and vulnerabilities. Training can help avoid risk.
- **Continuous monitoring**—Monitor the organization's policies and controls for continued effectiveness. Update policies and controls that aren't effective.

Figure 14-1 shows this basic RM process.

The RM process also can help an organization with its other contingency plans. The RM process and information learned during that process can form the basis for

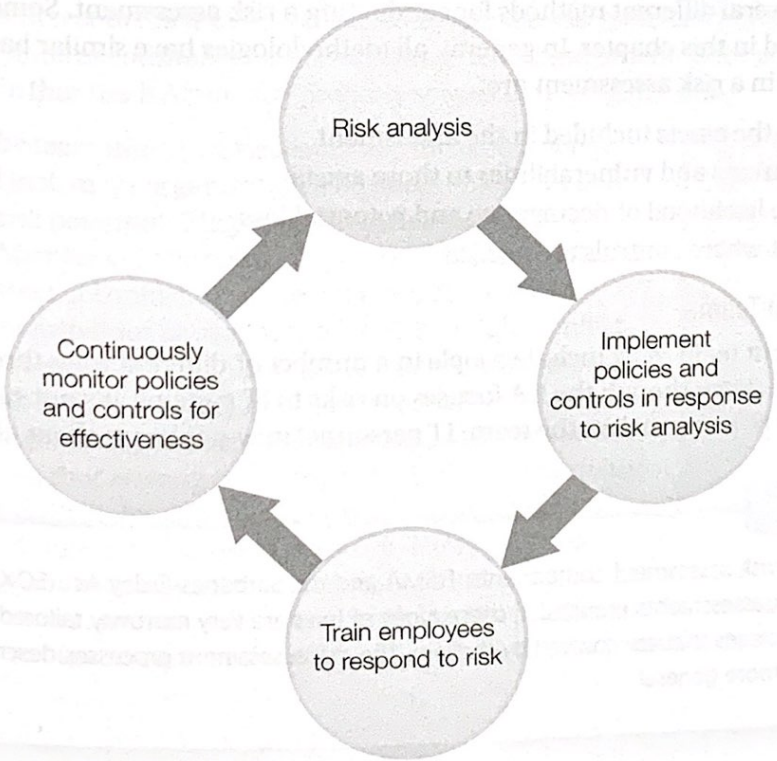

FIGURE 14-1 Risk management process.

an organization's contingency plans. In a basic sense, contingency plans are an organization's response to very narrow risk assessments. For example, business continuity and disaster recovery plans are detailed contingency plans that respond to the risks of natural or man-made disasters. These risks are discovered during the RM process.

Risk Assessment Process

A **risk assessment (RA)** identifies the threats and vulnerabilities to IT resources. It reviews the probability of those threats and vulnerabilities actually happening. This is called a **realized risk**. Then the RA reviews the potential harm from a realized risk. Finally, the RA identifies policies and controls that could respond to the potential risk. Risk response is the action taken by an organization to reduce realized risk to an acceptable level. The amount of risk that is left over after realized risk is reduced is called residual risk.

The RA is a tool in the risk management process. It provides executive management with the data that it needs to make smart decisions about information security controls.

Without an RA, the executive management team has no way to know if the controls and policies that it implements are needed or cost effective. For example, an organization could spend a lot of money implementing a control to respond to a perceived threat. Without an RA, the organization doesn't know whether the cost of the control is reasonable given the probability and potential of the underlying risk. The organization may be making a poor decision if it spends more money to respond to a risk than the actual cost of the realized risk itself. An RA helps the organization know where to spend money to respond to and mitigate information security risk.

There are several different methods for conducting a risk assessment. Some of them will be discussed in this chapter. In general, all methodologies have similar basic steps. The basic steps in a risk assessment are:

- Inventory the assets included in the assessment
- Identify threats and vulnerabilities to those assets
- Categorize likelihood of occurrence and potential loss
- Document where controls are needed

Risk Assessment Team

A risk assessment team must include people in a number of different roles throughout an organization. Even though the RA focuses on risks to IT systems, it's not sufficient to only include IT personnel on the team. IT personnel may not know about all of the

FYI

Many laws have risk assessment components. FISMA and the Sarbanes-Oxley Act (SOX) are two of them. The risk assessments required in these types of laws are very narrowly tailored to the systems and processes that are covered by the law. The risk assessment processes described in this chapter are more general.

organization's critical business processes. For the same reason, it's not enough to include only business personnel on an RA team. Business personnel may not appreciate the technology processes that support business operations. The RA team members should represent all areas involved in a business process workflow.

RA team members should include:

- **Business personnel**—These people are responsible for business process operations. They know the steps that must be completed in each business process. They also can describe how they use IT systems to accomplish their job duties.
- **IT personnel**—These people run the organization's various IT systems. This group also might include the IT system owner. These personnel understand how their IT systems work. They're responsible for maintaining those resources. They also would be responsible for implementing any changes to IT systems.
- **Information security managers**—These people run the organization's information security program. They have knowledge about information security threats and vulnerabilities. They also know how threats and vulnerabilities can be mitigated.
- **Human resources personnel**—These people understand how to deal with people issues. They can offer advice and input on how to address human-based threats and vulnerabilities. They also will be able to assist with awareness training after the RA is complete.
- **Executive management**—These people make sure that the RA team has the support and resources that it needs to complete the assessment. They can hold business units accountable for participating in the assessment. They also can make sure that the RA remains properly scoped to its original goal.

Members of the team must be objective. They must be willing to take a hard look at an organization's business objectives and processes for risk potential. They must do the same thing for IT resources. Members of the team must avoid **conflicts of interest**. In the RA context, a conflict of interest is a situation where a member's responsibilities as part of the RA team might conflict with their job responsibilities. A conflict of interest prevents the team member from meeting their RA responsibilities.

In many cases, team members are picked because they work in a certain area that is being reviewed as part of an assessment. These members are helpful in making sure that the whole RA team understands how these areas work. However, these members may not be as objective in assessing vulnerabilities, threats, and risks in their areas. These team members will have to work hard to be objective throughout the whole RA process. An organization must remove them from the RA team if they're unable to be objective.

> **NOTE**
> In general, a conflict of interest is any situation where a person's private interests and professional obligations collide. In these situations, independent observers might question whether a person's private interests improperly influenced his or her professional decisions.

Some organizations hire consultants to help with an RA. Consultants aren't employees of the organization. They help in a risk analysis by bringing objectivity to the team. They also help by bringing risk analysis expertise to the process. They can hurt the team if they don't quickly learn how an organization's business processes and IT systems work. They also can hurt the team if they aren't respectful of an organization's culture.

An RA team is responsible for collecting information about assets and risks. It's also responsible for reporting the assessment results to an organization's executive management.

Identifying Assets, Vulnerabilities, and Threats

A risk assessment must be narrowly scoped. Otherwise, it can grow too large for the RA team to manage. If it's too large, the team will have a hard time determining the risks that the organization must address. A good practice is to conduct a focused RA. Rather than doing an assessment of the organization's whole IT infrastructure, the RA team should review one infrastructure component at a time. The team can better manage the RA and produce better overall results.

> **NOTE**
>
> Many RA projects have a project manager. This person is in charge of making sure that the RA moves forward in a timely manner. The project manager isn't necessarily a participant on the RA team. Instead, this person helps the RA team meet their project goals and final deliverables.

The scope of the RA must be clearly stated in an RA project plan. RA team members can refer back to the project plan anytime they're asked to consider a new system or process as part of the RA. If the new system or process isn't within the scope of the project plan, then the team should not consider it. Executive management must approve the project plan before the RA team begins to work.

After the organization approves the scope of the RA, the RA team must identify all the assets that are within the scope of the assessment. This inventory must include the IT resources. It also must include a listing of the personnel who run business processes that are in the scope of the RA. The inventory also must contain the data included in those processes. For instance, an organization decides to conduct a risk assessment of its e-mail infrastructure. For this assessment, the RA team must consider not only the operation of the organization's e-mail servers, but also whether employees are sending confidential data via e-mail. Depending upon how the RA team defines the project, mobile devices that send and receive e-mail could be outside of the scope of the assessment.

The RA team should consider many assets for the inventory. They include:

- Personnel
- Data
- Hardware and software
- Physical facilities
- Business process workflows
- Current controls that help safeguard any assets

It's important to identify all assets that are associated with the RA project. The RA team must list assets that are critical to business functions. They also must identify threats and vulnerabilities to those assets. The RA team must identify vulnerabilities and threats in order to determine the risk to the organization's assets. For IT resources, the organization needs to know how vulnerabilities and threats might affect confidentiality, integrity, and availability.

A vulnerability is a weakness or flaw in an IT system. Information security is compromised when vulnerabilities are exploited. An exploit is a successful attack against a vulnerability. There are many different kinds of vulnerabilities, but they fall into four broad categories. They include people, process, facility, and technology vulnerabilities. Vulnerabilities can be design mistakes. They also can be configuration mistakes.

Threats are anything that can cause harm to an information system. They're successful exploits against vulnerabilities. A threat source carries out a threat or causes it to take place against a vulnerability. A threat source can be a person or a circumstance. Like vulnerabilities, threats also fall into four broad categories: human, natural, technology and operational, and physical and environmental. Threats can be deliberate or accidental.

Threat sources act upon vulnerabilities. Table 14-1 shows examples of vulnerabilities and threats.

TABLE 14-1 Examples of vulnerabilities and threats.

VULNERABILITY	THREAT SOURCE	THREAT
Data center has few physical security controls to prevent unauthorized access to data center hardware	Unauthorized users (e.g., criminals, terminated employees, curious employees)	Theft of data center hardware Theft of data on hardware Damage or destruction of hardware or data
Failure to remove user accounts in a timely manner when an employee leaves the organization	Terminated employees	Access to IT resources and theft of sensitive company data
No access controls for sensitive files stored on IT resources	Curious employees	Review of data without need to know Review of proprietary data Invasion of privacy of other employees and/or customers Unauthorized modification of data Theft of data

The RA team must put forth a good-faith effort to identify as many vulnerabilities and threats as possible. It can do this in a number of ways. For vulnerabilities that are internal to the organization, the team could interview employees in areas that are part of the assessment. The team also could send a questionnaire to a sampling of employees. The questionnaire could ask questions that are relevant to the assessment. The RA team also should review the organization's policies and procedures. The team could use automated scanning tools to look for vulnerabilities in its IT systems.

The RA team also must learn about vulnerabilities and threats that are external to the organization. They can review industry guides. They can use NIST guidance to learn more about information security issues and controls. The RA team can review information from the U.S. government and other industry experts. The SANS Institute has one of the largest collections of information security resources in the world. (See *www.sans.org* for more information.) The team also can review known system problems listed in the National Vulnerability Database (NVD). (See *http://nvd.nist.gov/home.cfm*.)

It's not possible for the RA team to identify every security vulnerability or threat. This is because technology changes so rapidly. For example, zero-day exploits are a type of exploit that's hard for an organization to prepare for. An organization can only prepare for zero-day exploits generally. It's hard to prepare for or assess specific zero-day exploits because they're exploited so quickly after they're discovered.

Likelihood and Potential Loss

A risk is the likelihood that a threat will exploit a vulnerability and cause harm. The RA team is responsible for determining how likely it is that identified risks will occur. The team also must determine the potential loss or harm that the organization could have if the risk is realized. Likelihood and potential loss can be stated as either a qualitative measure or a quantitative measure. It depends on the risk methodology that the RA team uses.

Quantitative Risk Analysis. **Quantitative risk analysis** attempts to use real numbers to calculate risk and potential loss. The organization assigns real numbers to the value of its assets. It also assigns real values to the cost of countermeasures and controls. One of the greatest advantages of a quantitative risk assessment is that it provides an objective assessment of cost. It helps an organization understand both the cost of risk and the cost of controls. This type of RA allows management to directly compare the cost and benefits of recommended controls. One of the greatest disadvantages of quantitative risk assessments is that they're very difficult to administer.

In a quantitative risk analysis, the RA team must assign a value amount to each of the organization's assets. A number of factors shape the cost of an asset, such as the cost of developing it and then the ongoing cost of maintaining it each year. The asset might have value that is hard to measure. For example, it might be hard to value the intellectual property that goes into creating and developing an asset. It's also hard to place a value on an asset that provides an organization with its competitive edge. While quantitative risk analysis does give an organization actual money amounts for its assets, there's a subjective element to these values that is hard to measure. This is particularly true if an asset has intangible value.

After the RA team determines the organization's vulnerabilities and threats, it must determine **exposure factor.** The exposure factor is the percentage of asset loss that is likely to be caused by an identified threat. For example, an RA team determines that an organization's data center is vulnerable to tornadoes. This is because the data center is located in the part of the midwestern United States known as "tornado alley."[3] The RA team estimates that if a tornado struck the data center, it would be destroyed. The data center's value would be reduced to zero. The exposure factor is 100 percent.

The RA team determines that the same data center is also subject to a threat of fire. The RA team determines that if a fire were to occur, only 50 percent of the data center would be destroyed. (This is because the data center has a fire suppression system). In this scenario, the exposure factor is 50 percent. Exposure factor also is called *likelihood.*

> **NOTE**
>
> Exposure factor is expressed as a percentage. It can range from zero to 100 percent. An exposure factor of 100 percent means that an asset is completely destroyed if a specified threat occurs.

Once the RA team determines the exposure factor, it must determine **single loss expectancy (SLE).** The SLE is the amount of money that an organization will lose if a risk is realized. It's the loss that an organization will suffer every time that risk occurs. SLE is often expressed as an equation:

SLE = Asset value × Exposure factor

In our data center example, suppose that the data center is worth $2 million. The SLE for the data center being struck by a tornado is $2 million. The SLE for the data center experiencing a fire is $1 million.

The RA team also must figure out how many times a specific risk might occur during a one-year timeframe. This is called the **annual rate of occurrence (ARO).** An RA team can use historical data to determine this number. It also can perform research to understand the likelihood of risks within a certain area. For instance, police departments can provide information about crime statistics. Insurance companies can provide information about how often organizations are likely to experience a serious fire or other devastating incident. ARO is expressed as a number. It can range from zero (a threat will never take place) to any number greater than zero. A risk that will happen only once a year has an ARO of one.

In our example, the RA team estimates that a fire might strike its data center once every 20 years. The ARO for that event is 1/20, or .05. The RA team estimates that a tornado might strike its data center once every 15 years. The ARO for that event is 1/15 or .067.

ARO is used to calculate **annualized loss expectancy (ALE).** The ALE is the amount of loss that an organization can expect to have each year due to a particular risk. In a quantitative risk analysis, the ALE is the end-result number. An organization can use this number to determine how much money it should invest each year in controls to mitigate a particular risk. As a matter of good business practice, an organization shouldn't spend more than the ALE amount each year on particular risks. ALE is often expressed as an equation:

ALE = SLE × ARO

In our example, the ALE for the tornado risk is $134,000 ($2 million × .067). The ALE tells the organization that it can afford to spend up to $134,000 per year to mitigate the potential loss caused by a tornado. The ALE for the fire risk is $50,000 ($1 million × .05). The organization can spend up to $50,000 per year to mitigate potential fire loss. Spending more on controls than the ALE amount might be unwise from a business standpoint. This is because the cost of the controls is then more than the cost of the risk. It doesn't make sense to spend more money on controls than the cost of the risk.

Organizations also can use these formulas to determine the value of a control. The value of a safeguard to the organization can be determined by comparing the ALE before implementing a control to the ALE after implementing the control. This also can be expressed as an equation:

Value of control = (ALE before safeguard) − (ALE after safeguard) − Cost of control

These equations are used in quantitative risk analysis. Many vendors have created software products that can help RA teams perform these equations.

Organizations often choose to do a quantitative RA because it computes risk in terms of money value. Sometimes executive management finds this information very helpful. It gives them concrete values upon which to measure the costs of risks and the effectiveness of controls. A quantitative risk analysis speaks the language of business: money.

Qualitative Risk Analysis. **Qualitative risk analysis** uses scenarios and ratings systems to calculate risk and potential harm. Qualitative risk analysis doesn't assign money value to assets and risk. Instead, it uses descriptive categories to express asset criticality, risk exposure (likelihood), and risk impact.

One of the greatest advantages of this method is that it's relatively easy to use. It doesn't require RA team members to have specialized financial knowledge. It also doesn't require them to deduce the costs of assets, controls, and potential harm. One of the biggest disadvantages of qualitative risk assessments is that they're very subjective. The opinions of members of the RA team can influence the results of the RA.

Qualitative risk analysis is scenario-based. In this type of assessment, the RA team considers a specific vulnerability or threat. The team then considers a scenario based on it. As team members walk through the scenario, they think about how the scenario will affect business operations and resources. The team members will use their own experiences to determine threat likelihood. They can also use industry resources.

The RA team will create likelihood and impact categories. These categories help them classify and compare different risks. These categories are usually based on either a numeric rating system (scale of 1 to 10) or a low-medium-high standard.

Risk exposure categories classify the likelihood of certain risks. An RA team might use the following categories to determine risk exposure:

- **Low**—Events that are unlikely to happen within a year
- **Medium**—Events that are somewhat likely to happen within a year
- **High**—Events that are likely to happen within a year

Risk impact is determined the same way. Impact categories classify the harm to an organization if a risk is realized. The RA team might use the following impact categories:

- **Low**—A realized risk will have little or no effect on the organization. The organization will experience light disruption to its business processes. It will incur only low costs related to lost productivity and data. It won't experience reputational loss.
- **Medium**—A realized risk will have a moderate effect on the organization. The organization may experience moderate disruption to its business processes. It will incur moderate costs related to loss productivity and data. Its reputational loss will be moderate.
- **High**—A realized risk will have a severe effect on the organization. The organization may experience severe disruption to its business processes. It will experience high costs related to loss productivity and data. Its reputational loss will be significant.

The RA team also will need a way to compare likelihood and impact levels to determine overall risk level. The RA team will use the overall risk level to determine where the organization should apply controls. An organization won't want to implement costly controls on a system that has low exposure and impact results. However, it will want to implement controls where there's a high likelihood rating and a high impact rating. It must create a method to analyze likelihood and impact rating levels. It uses this analysis to determine the priority of issues that the organization must address. Table 14-2 shows a generic risk level matrix.

Organizations use their risk level matrix to determine the priority of the risks that they must address. Table 14-3 shows what the results of a qualitative risk analysis might look like. It uses the vulnerabilities and threats from Table 14-1. In this example, the RA team would first recommend that the organization implement controls to remove user accounts for terminated employees in a timely manner. This is because the risk level for that vulnerability is higher than the others listed in the table.

The problem with qualitative risk assessments is that the organization has no way to determine the amount of money to spend on controls. It also has no way of knowing if it's spending too much on a control relative to the actual loss that it could have due to a realized risk. For some organizations, this is a significant failure of a qualitative risk analysis. This is one reason why an organization might prefer to complete a quantitative risk analysis. Table 14-4 compares the features of qualitative and quantitative risk assessments.

TABLE 14-2 Risk level matrix.

		LIKELIHOOD		
	RATING	**Low**	**Medium**	**High**
IMPACT	**Low**	Low	Low–Medium	Medium–High
	Medium	Low–Medium	Medium	Medium–High
	High	Medium–High	Medium–High	High

TABLE 14-3 Risk level outcomes.

VULNERABILITY	THREAT SOURCE	THREAT	THREAT LIKELIHOOD	THREAT IMPACT	RISK LEVEL
Data center has few physical security controls to prevent unauthorized access to data center hardware	Unauthorized users	Theft of data center hardware	Medium	High	Medium–high
Failure to remove user accounts in a timely manner when an employee leaves the organization	Disgruntled terminated employees	Theft of sensitive company data	High	High	High
No access controls for sensitive files stored on IT resources	Curious employees	Review of data without need to know	Medium	Low	Low–medium

TABLE 14-4 Qualitative and quantitative risk assessment comparison.

	QUALITATIVE RISK ASSESSMENT	QUANTITATIVE RISK ASSESSMENTS
Positive Aspects	Easy to administer Doesn't require formal knowledge to administer Calculations are simple Scope of the assessment can be changed easily if necessary	Measures the money cost of a risk Very objective, can be used to make cost-benefit decisions Easy for executives and financial managers to understand
Negative Aspects	No measure of money cost of risk Very subjective, can't be used to make cost-benefit decisions	Hard to administer Requires formal knowledge to administer Calculations are complex Must put effort into valuing assets (and that value may be hard to calculate) Very difficult to change the scope of the RA

Document Needed Controls

The final part of an RA is to document where security controls are needed. The RA team does this by reviewing the results of its assessment. In a quantitative risk analysis, the RA team has a list of risks that it can prioritize by ALE and ARO. The RA team should review risks with high ALEs and high AROs to recommend controls.

Risk Assessment Methods

There are a number of different RA processes to choose from. Some vendors offer software that organizations can use to complete RAs. This sidebar describes some of the RA methods you might hear about in the information security profession.

NIST "SP 800-30 (Rev. 1): Guide for Conducting Risk Assessments" is a qualitative RA method. Any type of organization can use it. It describes a number of tasks to be completed in the RA process. You can view this NIST guide at *http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf.*

Britain's Central Computer and Telecommunications Agency (CCTA) created CRAMM (CCTA Risk Analysis and Management Method) in 1987. CRAMM has been commercialized and is sold as a software program. It's a qualitative RA method. It relies on international standards for its RA methodology. CRAMM advertises itself as an ISO/IEC 27001-compliant RA methodology. You can learn more about it at *http://rm-inv.enisa.europa.eu/methods/m_cramm.html.*

Another framework developed in Britain is the M_o_R framework. This framework describes the processes that need to be put in place to implement risk management. You can read more at *http://www.mor-officialsite.com/.*

Carnegie Mellon University's Software Engineering Institute developed OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation). OCTAVE is a qualitative RA process. It's designed to be self-directed. It allows organizations to align security operations to business strategy. You can learn about OCTAVE at *http://www.cert.org/resilience/products-services/octave/index.cfm.*

Thomas Peltier developed the Facilitated Risk Analysis and Assessment Process (FRAAP) in 1993. It's a three-step process that helps organizations identify risks and threats. It's based on ISO/IEC 27001.

The Microsoft Corporation describes a risk analysis process that contains both quantitative and qualitative elements in its "Security Risk Management Guide." It created this guide in 2006. You can learn about it at *http://technet.microsoft.com/en-us/library/cc163143.aspx.*

Through a qualitative risk analysis process, the RA team has a list of risks organized by potential harm. The risk level is based on the classifications that the RA team used for risk likelihood and impact. The RA team should look at high and medium risk levels to recommend controls.

An RA team needs to suggest controls to executive management. It must identify risks that are high priorities. It also should identify controls that can mitigate or eliminate those risks. If the RA team performed a quantitative analysis, it should include a cost-benefit analysis for each of its suggested controls. If the team performed a qualitative analysis, it should show how the suggested control affects the overall risk level for a specific threat or vulnerability.

Executive management reviews the RA team's report as it makes business decisions. It will use the report as part of its information security governance (ISG) activities. The organization must make sure that its response to the risks identified in the RA team's report supports its business objectives. The executive management team shows due care and due diligence when it responds to issues that are raised in a risk assessment.

Risk Response

> **NOTE**
> Controls reduce the harm posed by vulnerabilities or threats. They may eliminate or reduce risk of harm. They're also called safeguards or countermeasures.

Executive management must decide how an organization responds to risk. Risk response is the actions taken by executive management to reduce risk to an acceptable level. Executive management must apply the most appropriate controls to decrease its risk. The organization's risk response must be cost-effective. It also should have a limited impact on the organization's business.

Executive management must prioritize how it will respond to risks. It should respond quickly to the risks that have the greatest potential to harm business goals. Executive management must assess each risk individually. An organization won't handle all risks in the same manner. Executive management can use a number of approaches to respond to risk. They are:

- **Risk avoidance**—The organization applies controls or takes other action to completely avoid a particular risk. This strategy removes all risk caused by a particular vulnerability or threat. For example, executive management could decide that the risks posed by a function in an IT system outweigh the benefit of that function. It could instruct system owners to disable that function in order to avoid the risk. The risk caused by that function is completely eliminated.
- **Risk mitigation**—The organization applies controls or takes other action to reduce a particular risk. This strategy doesn't eliminate all harm that could be caused by that risk. Instead, it reduces the risk to an acceptable level. The risk that is left over is called residual risk. For example, executive management could decide that the risks posted by a function in an IT system could be lessened if access to that function was limited. The organization could use access controls to limit access to that function to only a few trusted employees. The risk caused by that function is reduced because fewer people have access to it.
- **Risk transfer**—The organization takes no action against a particular risk. Instead, it passes its risk to another entity. The other entity bears the risk of loss. Usually an organization transfers risk of loss to an insurance company. It can purchase a cyberliability insurance policy to insure against specific risks. For example, the organization could purchase insurance to cover losses due to unauthorized access to IT systems. These types of policies are called technology errors and omissions policies.

- **Risk acceptance**—The organization takes no action against the potential risk. It makes an intentional decision to do nothing. Executive management may choose this strategy if the cost of the risk is less than the cost to avoid, mitigate, or transfer the risk.

> **NOTE**
> Executive management can implement administrative, technical, or physical controls to avoid or mitigate risk.

Training Employees

The RM process includes educating employees about risk and how to avoid it. Often employees contribute to an organization's risk. They engage in behavior that puts the organization and its data at risk. Risk can be reduced if employees understand why certain behaviors are not allowed or acceptable. The organization can use policies to help direct employee behavior. Risk is reduced when employees behave according to company policy.

Security awareness training is an important part of any information security program. It's required to keep employees aware of their security responsibilities. An organization must include security education in its day-to-day routine.

Continuous Monitoring

Organizations must always monitor their information security risk. They also must monitor the controls that they put in place to respond to that risk. This is an ongoing process. Since technology changes quickly, executive management must be diligent reviewing its risk and response to that risk.

Many laws require organizations to continually assess their risk. For example, the Gramm-Leach-Bliley (GLBA) Safeguards Rule requires covered financial institutions to conduct RAs to identify risks to customer information as part of their information security program. Covered financial institutions must apply controls to respond to identified risks. They also must assess their current controls to make sure that they're effective. Financial institutions also must review their information security programs on a regular basis.

The Health Insurance Portability and Accountability Act (HIPAA), FISMA, and SOX also require covered organizations to conduct regular risk assessments.

> **NOTE**
> Many U.S. laws require federal agencies and other organizations to engage in regular risk assessment and management activities.

Three Types of Contingency Planning

One of the most important things about conducting an RA is that it forces an organization to think "What if?" An organization must plan for the risks that it accepts or can't avoid, mitigate, or transfer. It creates contingency plans to limit the financial loss it might experience due to an adverse event. Contingency plans help to minimize the length of time that services and processes are interrupted. They also help minimize customer impact due to a serious event.

The three main types of contingency plans are:

- Incident response plans
- Disaster recovery plans
- Business continuity plans

An organization creates these plans to respond to events that might negatively affect IT resources and business processes. Keep in mind that this chapter discusses these types of contingency plans from an information security perspective. These plans can have a much larger scope than just IT.

It's also important to remember that the most important goal of any type of contingency plan is to preserve human life. Other goals are secondary.

Incident Response Planning

> **NOTE**
> Incident response is a reactive term. It describes how an organization responds to an incident. Incident handling is a proactive term. It describes how an organization manages an incident. Organizations may use both terms interchangeably.

An organization uses its **incident response (IR)** process to react to attacks against its IT infrastructure. Having an incident response process is important because it helps make sure that an organization can recover from security incidents. Organizations that are able to recover quickly from incidents are more likely to be able to continue business operations. Incident response is also called incident handling.

Incident response describes how an organization:

- Detects information security incidents
- Determines the cause of the incident
- Mitigates the damage caused by the incident
- Recovers from the incident

An **incident** is any event that involves the organization's equipment, data, or other resources. An incident must adversely affect the confidentiality, integrity, and/or availability of the organization's data and IT systems. The intent of the threat source does not matter. An incident includes malicious attacks. It also includes the harmful acts of well-meaning employees. Incidents are usually violations of the organization's policies, accepted security practices, or the law.

Organizations encounter a number of information security threats each day. Each one of these can be an incident if it adversely affects the security of an organization's resources and data. You hear about these incidents in the media nearly every day. For example, in 2007 student hackers installed keystroke loggers on some computers in a university registrar's office. They gathered username and password information from the keystroke loggers. They then used that information to change grades in the university's computer system. The university discovered the hack through logging and audit review security measures.[4]

When the university discovered that records had been improperly modified, it would have declared an information security incident. A team would be formed to respond to the system intrusion. That team would have been responsible for determining the damage

that the incident caused. They also would find the source of the incident. They also would mitigate the damage caused by the incident and recover the computer systems affected by the incident.

The student hackers' conduct was most likely a violation of the university's acceptable use policy (AUP). Many organizations have these policies to define acceptable behaviors for the use of its IT resources. In this case, the student hackers violated policies and also broke the law. One of the students was sentenced to seven years in prison for violating the federal Computer Fraud and Abuse Act.[5]

Based upon what we know about this case from publicly available information, we know that the university involved did the following:

- **Detecting the incident**—The university used normal information security practices such as logging and audit review to discover that information in an IT system had been changed.
- **Determining the cause of the incident**—These same measures, and most likely computer forensic analysis, helped the university discover that the cause of the incident was student hackers who had installed a keystroke logger.
- **Mitigating the damage caused by the incident**—The university was able to restore the changed data because it made regular data backups.
- **Recovering from the incident**—The university restored data and, likely, hardened its systems to prevent this type of incident from happening again in the future.

One of the most important parts of IR is documentation. An organization must document every incident that it encounters. That way it can refer to its documentation if it encounters a similar incident in the future.

Incident Response Team

The IR team is responsible for creating the organization's IR policy and plan. This team has a different focus from the operational information security employees that will follow the IR plan and respond to incidents. Often the IR team will include many information security team members. Like an RA team, the IR planning team must include advisors from a number of departments across the organization.

The IR team will help draft the initial IR policy. It also will create the plans that define the organization's IR structure. Even though IR is often an information security responsibility, other departments may need to participate in an IR process. The information security team may have to interact with many of these departments when handling an incident.

For instance, the operational team will be involved if there's an incident involving a physical trespasser to the organization's data center. The operational team most likely will involve the organization's physical security personnel in handling the incident. This is because the incident included actual trespass onto the organization's property. IT personnel will be involved to see if any equipment has been stolen. Internal audit

and legal counsel will be involved if any equipment or data are stolen. Human resources (HR) personnel could become involved if it appears that the trespasser is a former employee. Marketing and communications personnel could become involved if the facts of the incident trigger the laws that require security breach notification.

The IR team is responsible for making sure that the procedures are in place to help all of these different departments work together. They must be able to respond quickly and efficiently in the event of an incident. The operational information security team needs to know whom they can contact in each department for help. The IR team's planning puts this structure in place.

The IR planning team should include information security and IT representatives. It also should have members from physical security, HR, and internal audit. Legal counsel should be included on the team to address any legal or regulatory issues.

The IR team will help the organization create its IR policy. These policies are very specific to an organization's structure and culture. Like all information security policies, the IR policy is a statement of executive management's commitment to the IR process. The policy should state the purpose and goals of IR. It also should define what an incident is. The policy must set forth, at a high level, the organization's operational IR structure. It should define the roles and responsibilities within this structure. The IR policy also can contain information about how to measure the effectiveness of the IR process.

IR Plan Process

Once the organization approves its IR policy, the IR team can continue to define the IR plan. The IR policy contains the overall IR goals. The IR plan contains the procedural elements that are necessary to meet those policy goals. Operational information security teams will follow the IR plan to fulfill their IR job duties. In this section, the term *incident handlers* will be used to refer to the operational teams that respond to an incident.

An IR plan is specific to a particular organization. However, most IR plans have five basic parts. They are:

- Incident triage
- Investigation
- Containment or mitigation
- Recovery
- Review

> **NOTE**
>
> The word *triage* is most commonly associated with the medical profession. It's the process of sorting and prioritizing patient care based on the severity of a patient's condition.

The triage phase is the first phase in the IR process. This is the phase where a potential incident is initially assessed. It's at this point that the primary handler will verify whether an adverse event meets the definition of an incident. The primary handler is the person who is in charge of coordinating an organization's response to an information security incident. This person is often a member of the organization's information security team.

The Operational Incident Response Team

An organization often plans and coordinates IR at a high level within an organization. A cross-functional team, called the IR team, puts the IR plan into place. The responsibility for the daily operations of the IR plan often falls to an organization's information security team and other IT personnel.

There are a number of different roles involved incident response. These roles are reviewed briefly here. You may find these terms used in resources describing how to plan and implement an IR program:

- **Victim**—The person or resources that are targeted in an incident. The victim is often the organization and its IT resources and data.
- **Attacker**—The person or mechanism that caused the incident.
- **Incident reporter**—The first person or mechanism that reports an incident. The incident reporter doesn't have to be a person. An automated intrusion detection system (IDS) can be an incident reporter. A person who notices an unusual incident and reports it is also an incident reporter.
- **Primary handler**—The person who is in charge of coordinating the response to a particular incident. This person is responsible for making sure that the IR process is documented. Often this person is a member of the organization's information security department. If an organization doesn't have a dedicated information security department, this is the person with information security job duties.
- **Secondary handlers**—These are the personnel involved in investigating, responding, and recovering from an incident. Secondary handlers include technicians, analysts, and operational staff who take part in handling the incident. Legal counsel and an organization's internal auditors also can be secondary handlers. The type of secondary handlers involved in IR depends on the nature of the incident.

Not all events are incidents. The event must have an adverse effect on the confidentiality, integrity, and/or availability of an organization's IT resources or data. For example, it might be an incident if an employee mistakenly deletes a critical file needed for processing the organization's weekly payroll. The employee has compromised the availability of needed data. It might not be an incident if the employee mistakenly deleted a non-critical file. (Although, it certainly would be a business process issue.) The primary handler decides whether a reported event is actually an incident.

If an event is an incident, the primary handler must classify it. Incidents can be classified in a number of ways. They can be sorted based on threat source. An example would be whether the incident occurred due to an internal threat or an external threat. Incidents also can be sorted based upon the type of vulnerability or threat that's exploited.

Organizations may use any method for sorting incidents. Some organizations use guidance prepared by NIST.[6] The United States Computer Emergency Readiness Team (US-CERT) uses incident response categories that are based on NIST guidance. The US-CERT is the federal government's IR center. All federal agencies must report information security incidents to the US-CERT. The US-CERT's incident categories are:

- **Category 1: Unauthorized Access**—Unauthorized access is technical or physical access to an IT system without permission. An agency must report these incidents even if data isn't compromised.
- **Category 2: Denial of Service (DoS)**—Any event that prevents the normal operation of IT resources such that use of those resources is harmed.
- **Category 3: Malicious Code**—Any event that involves the use of malicious code to successfully infect, breach, or compromise IT resources. These events include viruses, worms, and Trojan horses.
- **Category 4: Improper Use**—Any event that is a violation of the agency's AUP or other related policies.
- **Category 5: Scans, Probes, and Attempted Access**—Any event where an IT resource is scanned or probed in an attempt to access or identify the agency's IT systems.
- **Category 6: Investigation**—This category is for unusual events that don't fall into one of the other categories. These incidents require more review because they're odd or potentially harmful.

The classification of an incident may change as the incident is investigated. This is because the organization learns more about the incident as the investigation progresses. Incidents also should be sorted based upon potential severity. Severity is assessed based upon the perceived level of impact to the confidentiality, integrity, and availability of an organization's IT resources or data. The severity of an incident also may change. Organizations often classify severity on a low-medium-high scale. An organization might classify severity as follows:

- **Low**—The adverse affect on the confidentiality, integrity, or availability of the organization's data or IT resources is limited. A low-impact event causes little or no damage.
- **Medium**—The adverse affect on the confidentiality, integrity, or availability of the organization's data or IT resources is moderate. A medium-impact event results in significant damage to assets.
- **High**—The adverse affect on the confidentiality, integrity, or availability of the organization's data or IT resources is severe. A high-impact event results in major damage to assets.

Classifying the nature and severity of an incident helps incident handlers know which incidents require priority handling. If the incident handlers must respond to multiple incidents, classification promotes efficiency. It also makes sure organizations respond to incidents according to the incident's potential to hurt the organization. Classification also helps incident handlers know which incidents to escalate to management.

Investigation is the second phase in the IR process. During this phase, the incident handlers learn about the incident and its source. They learn more about the impact that the incident is having on the organization. The incident handlers must find all the resources that are affected by the incident. They also must contact other areas as needed to fully understand the scope of the incident. The IR policy and plan let the incident handlers know who they must contact.

The incident handlers must keep management informed of their IR activities. It's important that management be informed in case there are any regulatory requirements that need to be considered as the organization responds to the incident. For example, if an incident involves the disclosure of the protected health information of more than 500 people, HIPAA requires that the organization notify the Department of Health and Human Services about the disclosure.[7] Executive management must make the final decisions about contacting third parties or the media.

It's also important during the investigation phase for the incident handlers to follow the organization's own internal policies. If handlers are investigating an incident that might be a crime, it's important that the team follow good evidentiary practices. If an incident appears to be a crime, the incident handlers must contact law enforcement according to the terms of the IR plan.

The organization must begin the containment phase almost as soon as an incident is reported. During this step, incident handlers must take steps to limit the damage caused by the incident. They will use different methods to contain an incident depending upon its nature. For example, if the incident is a self-propagating virus, incident handlers may remove an infected system from the organization's network. If the incident is a particularly authentic-looking phishing e-mail, incident handlers might issue an alert or other notification to the organization's employees. The alert would tell the employees not to respond to a phishing e-mail. Incident handlers might use a number of different tactics to mitigate an incident.

The organization repairs and recovers its IT resources and data during the recovery phase. The IT resources and data should be repaired in such a way that they're not vulnerable to the same type of incident again. This is called hardening. In addition to hardening damaged IT resources, the organization must harden resources that are similar to the damaged resources. This makes sure that similar resources aren't harmed by similar incidents.

After an IT resource is repaired, it should be tested for any additional vulnerabilities or weaknesses before it's put back into production. The recovery and repair method will depend upon the nature of the incident.

All stages in the IR process must be fully documented. The primary incident handler is responsible for making sure that each step in the process has been fully documented. This is important because the IR planning team can review the notes from the incident handlers to determine whether the IR plan worked as intended. During the review stage, both the incident handlers and the IR planning team can review the documentation to learn:

> **NOTE**
>
> As a general information security best practice, a repaired IT resource should not be tested by the same person that repairs or recovers it. This is a separation of duties best practice to make sure that vulnerabilities and weaknesses aren't overlooked.

- Dollar amount spent in handling the incident
- Dollar amount spent to prevent similar incidents in the future
- Loss of staff time in handling the incident
- How the response to the current incident compares to similar incidents in the past
- Recommendations on policy and procedural changes as a result of lessons learned from the incident

The review phase is often overlooked. This is because it's easy for an organization's employees to go back to their normal operational duties after an incident. The IR planning team must make sure that the review process is formally required by policy.

The IR process is shown in Figure 14-2. As described in this section, many of the stages overlap with one another. The lines between each stage can be indistinct at times.

Disaster Recovery and Business Continuity Planning

A 2013 AT&T study found that 87 percent of the organizations surveyed said that they had a business continuity plan in place. Sixty-three percent of the organizations said that security breaches were their most important security concern in 2013.[8]

> **NOTE**
> It helps to think of an incident as an event that an organization can deal with during its normal operations. A disaster is an event that completely disrupts those normal operations.

Disaster recovery and business continuity plans help an organization respond to a disaster. A **disaster** is a sudden, unplanned event. Disasters negatively affect the organization's critical business functions for an unknown period. The difference between a disaster and an incident is subtle. A disaster severely affects the organization's infrastructure. It interrupts critical business functions. An incident tends to refer to service failures that affect the confidentiality, integrity, and/or availability of the organization's data and IT systems. An incident may become a disaster in some situations.

Examples of disasters include natural threats and deliberate, human-made threats. Natural threats are uncontrollable events. They include earthquakes, fires, and flood. Human-made threats include sabotage and terrorist activities. Threats that can evolve into a disaster also include equipment, electrical, and communications infrastructure failure. Disasters aren't predictable. Organizations can't control these types of threats. They can take measures to try to limit the damage caused by these events. Organizations also can make plans to respond to these types of events.

FIGURE 14-2
Incident response plan phases.

Disaster recovery (DR) plans focus on how an organization recovers its IT systems after a disaster. These plans focus on IT systems only. They're the organization's immediate response to restoring critical IT resources after a devastating disaster or event. DR is largely a function of IT and is part of a larger business continuity plan.

Business continuity (BC) plans focus on how the organization continues its business during and after a disaster. These plans tend to be more comprehensive. They cover all parts of a business, not just IT systems. These plans address the period between the disaster and a return to normal operations.

The formal distinction between DR and BC plans is eroding. This is because organizations rely on IT systems for many of their critical functions. Some organizations build their whole business model on their IT systems. Today, most organizations consider DR and BC to be the same thing. You may find the terms used interchangeably. Most references recognize that a DR plan must be part of a comprehensive BC plan.

In this chapter, DR and BC plans will be discussed together, unless it's necessary to differentiate between the two. If it's necessary to differentiate between the two types of plans, the distinction will be made clear in the text.

DR/BC Team

The DR/BC team is responsible for creating an organization's DR/BC policy and plans. Like the RA and IR teams, this team must include members from many areas of the organization. The people who are going to be responsible for carrying out the plan in the event of a disaster also should be included on the team.

A DR/BC plan is an organization-wide plan. Specific departments may have secondary DR/BC plans. These department-level plans must be consistent with the organization's DR/BC plan. An organization's overall DR/BC plan has a number of goals:

- Ensure that the organization's employees are safe
- Minimize the organization's amount of loss
- Recover critical business systems and infrastructure within a certain period
- Resume critical business operations within a certain period
- Repair or replace damaged facilities
- Return to normal operations

The DR/BC team will help the organization create its DR/BC policy. These policies are very specific to an organization's structure and culture.

A DR/BC team must include members from IT, HR, executive management, physical security, and legal counsel. It must include the managers or owners of critical business processes. The team also should include the employees, and their backups, who will be in charge of directing the organization's activities in a real disaster. It's important that backup personnel be included in case the people with primary DR/BC responsibilities aren't available in a disaster.