Chapter 9

Continued Assessment of Needs, Threats, and Solutions

Chapter Objectives

After reading this chapter and completing the exercises, you will be able to do the following:

- Determine the lessons that were learned during the test disaster recovery.
- Decide how to overcome the threats that were uncovered.
- Use SWOT (strengths, weaknesses, opportunities, threats) analysis as an additional method of determining threats.
- Plan for eliminating threats going forward.

Introduction

The test is done; now it is time to determine what was learned in the test, and what can be taken from the experience and used so that the same issues (and there will be issues) aren't repeated the next time.

What to Do After the Disaster Recovery Test

After the test, it is important not only to step back and regroup, but also to meet as a team and discuss the lessons learned. This is not the time to assess or assign blame. This is the time to simply determine what was learned and what can be done better going forward. It is important that the team meet, before a long enough period of

downtime that the they begin to forget, at least once (typically more than once) to discuss what was done well, what was done poorly, and what was learned.

Extensive notes should be taken during these meetings so everyone's perception of the lessons learned is captured. In the aftermath of the recovery or the recovery test, some thoughts may not seem important or relevant, but after a period of time they may appear to be more relevant.

What Was Learned?

The recovery test is over. The team has learned that either the organization can or cannot recover, and that is all that was learned, right?

Not exactly. There are many things that may have been learned. For example, the organization may be conducting backups at the wrong time. Missing critical information in a given time period may be captured if the time is changed by even a small fraction.

Perhaps the organization cannot recover because the tapes being used are incompatible with the hardware on which it is trying to recover. This can mean simply changing the kinds of tapes used, or changing the kind of hardware used for recovery. This depends largely on the recovery site and the recovery solution that may be under contract.

Perhaps there were missed applications, files, or databases that should have been backed up in order to fully recover the organization's systems. These missed applications are often the ones that are sitting on someone's PC that they never think of even being there until they need it; they are the only ones who use it, and they are the only ones who will notice that it is missing. This information needs to be captured so changes in strategy can be made.

There are often applications that people forget that they rely upon for their day-to-day jobs. For example, a reporting database, written in Access and residing on a user's computer but shared among all of the people in the department, might not make the list of critical applications. This application would likely be homegrown and might make use of information gathered from multiple databases and massaged into a format that is easier for the end users to work with. Without this application, it would take several times longer for the users to perform their daily reporting and decision-making duties. The users likely would take this application for granted (after all, it was likely not purchased but built internally by a member of the group rather than a member of the IT development staff), the computer on which it resides may not make it into the backup schedule, and the information contained could be lost.

Depending on your backup strategy and how much information and backup software you keep with each copy of your recovery backups, you may find that, when you attempt to recover, you are unable to because you have not taken into account that the software may not be loaded on the computer on which you recover. Or you may find that the version of the software that you have available is not the version that is required to read the files that you have

restored. You may have upgraded your installation, but not the backup files of the software to reflect this upgrade.

Any number of things can be uncovered or unrecoverable during either the recovery testing or the recovery proper. It is just as important, however, to determine things that may have gone particularly well. If you have gone through the recovery process a number of times, you may have a general feel for how long each piece of the recovery process will take to accomplish. Perhaps someone in the organization came up with a new way to do the backups so that the recovery could be accomplished in half the time. The time savings shortened the timelines on the critical path and all the systems dependent on the shortened recovery were likewise finished early. This was brought out as a lesson leaned from the recovery at the next postmortem debriefing.

Although it is important not to assign blame for things that went badly, it is often beneficial to morale to give kudos for things that went particularly well. It is often to the team's benefit to praise even minor achievements. Although it is important to not appear to be insincere, it is also important to provide positive reinforcement.

IN PRACTICE: Lessons Learned

Radical Risk Management, a Miami, Florida company, had been holding recovery tests every January in Denver, Colorado for several years. The recovery team flew in to the recovery area, as was their standard practice. They requested that the company that stored their tapes send them to the recovery area so they were there when the team arrived, as was their standard practice.

The team arrived on time, but the tapes were running a bit behind schedule. When the tapes finally arrived, there was ice inside the tapes and water inside the tape case. Apparently there were problems during shipping (presumably only during shipping) and the tapes appeared to be unusable.

The team, in the spirit of disaster recovery, made the decision to attempt to recover anyway and see what they could salvage and what they could learn. After filing a claim with the tape storage company, they spent the entire first day of their recovery effort trying to dry out the tapes and restore them to a condition where they could be read from and written to.

Although they were unable—due in part to lack of time and in part to the degraded condition of the tapes—to completely recover, they were able to take back a new set of lessons learned on new and innovative ways to dry out wet tapes, which they shared not only with their organization, but also with their trading partners too.

During these sessions, it is important to remember to stay on task as much as possible. It is also important that the meetings not become a place to judge either processes or people but to find facts. If the meetings become a place where blame is assigned, people will become defensive and little will be learned. Remember, the team is just that—a team.

What Will Be Done Differently

Once the fact-finding meetings are complete, and this has to be decided by each individual organization, it is time to make decisions on what the organization, and the team, will do differently during the next recovery (whether it is a test or an actual recovery).

There are many places where things are likely to change, depending on what was uncovered in the meetings. There may simply need to be alterations made in the documentation, or there may need to be changes made in the overall process. Some will be small and simple to implement; others will be broader in scope and scale and will take cross-functional effort.

These meetings on changes to make should occur as soon after the recovery as possible. This is for two reasons: first, to make sure that you don't lose the momentum gained at the recovery and immediately following; and second, to allow as much time for the implementation of the changes as possible before the next recovery effort.

Changes should include any measures that can be added to make security tighter at the next event. These incremental security changes will be simpler to implement than if several iterations pass with no changes and then it becomes apparent that the changes are necessary. Because disasters deal almost entirely with security or lack of adequate security, it is important that security stay uppermost in everyone's mind.

The recovery planning process often seems to become rote and the real purpose lost in the process. People forget that the purpose is to recover the organization and get lost in the details of getting through the meetings and the recoveries. It likely will not be possible to get the people really excited about the prospect, but keeping a level of awareness and excitement is something that needs to be a goal of the team.

It is important that every item that ends up in the "what we learned" category be at least examined in the "what we will do differently" meetings. Even if all that the group does is state the situation and open it for discussion, at least the issue will have been aired and addressed as relevant or irrelevant and can then be put to rest. Not only does this allow all issues that arise to be addressed, it also allows the people involved to feel validated and their opinion valued. Remember, morale is as important in the recovery team as it is anywhere else.

Threat Determination in System

Based on the outcome of testing and the follow-up debriefing meetings, it is possible to determine what the uncovered risks are for the system. These threats include spoofing, tampering, repudiation, information disclosure, and denial of service, and level of privilege. It is important to note that in a static environment, for a given system design, the threats to that design don't really change. However, as the system evolves or as the systems surrounding it and feeding it evolve, so do the threats to the system. Organizations are, if nothing, continuously changing. This is becoming more and more the case. Look at *Who Moved My Cheese* and *Managing Chaos and Complexity*, two books about managing people and organizations in their times of change.

Ironically, once they have a disaster recovery plan, and even more so after they have started to test the plan, many organizations simply ignore that there may be emerging threats to their organization. They may feel that they have done due diligence and that is all they need to do. If they are lucky, they are right. It is more likely that something that they didn't notice sneaking up on them will cause issues later, or there will be an issue with the recovery site that they never considered. After all, it is a recovery site; they are supposed to have plugged all the holes, right? Maybe. Maybe not. Why take the chance?

Think about it. A new system may be created that feeds its information into a system that you are already backing up as a part of the recovery effort. The new system may or may not be included in the backups, but the connectivity between the new source system and the old target system may not be accounted for, and the dependencies between the two systems may not have made it into your documentation. This might mean that your recovery efforts are misplaced and that you are now recovering systems in the wrong order.

Or there may be holes in the security system that have been plugged in the primary location but that have not been plugged in the recovery site location. These security issues may include the screening of people who have the privilege to access the systems at the recovery site. You may find that information ends up being misappropriated by people that you are trusting at the recovery site or that people on the outside have gained access to your information in ways that you had not anticipated, but trusted the recovery site to have taken into consideration.

Threat Classification

In classifying the threats to each system and to each component of a system (including the human components), you can look at each kind of threat and determine what kinds of attacks associated with the threat could be launched at each component.

Each kind of threat attack can be further broken down as to whether they are mitigated threats or unmitigated threats. After recovery testing, the team should go back through the systems and look at the threats and determine which are mitigated (ones that have been taken into account in the recovery plan or removed from the equation) and which are unmitigated (those which still allow for vulnerabilities). Many unmitigated risks may go undiscovered until the tests have been conducted.

It is often beneficial to have someone on the team designated to attempt to find the holes in the system while the recovery is taking place. This is, in theory, the closest you are going to get to a replica of your production system, and allowing access to the system by someone who is determined to find the places where it may be broken is often a benefit. It may even be to an organization's benefit to schedule at least one test cycle that is dedicated to finding all the places where the system can be broken.

By throwing as many people at the problem as possible, it is possible to find many previously undetected holes. An organization may even consider hiring a consultant with experience in locating weaknesses in a system to point them out. Examples of some potential holes follow.

Spoofing Spoofing refers to one person or entity electronically masquerading as another by falsifying data, redirecting URLs, or redirecting messages to an alternative site. This masquerading can lead to an illegitimately gained competitive advantage by scavenging information from the business or to illegal activities by stealing information or money.

Spoofing can be the end result of many different kinds of attacks. Any attack that gains someone information can result in that information being used to spoof others into revealing even more information to the hacker or to trusting that person with business to which they should not be a party.

Spoofing can include phishing or Web page spoofing. In phishing, an organization's Web page or Web page design is copied (granted, easy enough to copy directly from the Internet, but why bother if you have access to the information directly from backup tapes, from recovery site servers, or from downed servers in a primary location left unguarded in an emergency situation), and unsuspecting victims provide personal information or make purchases of things that are not legitimate. This attack is often directed at banks and financial institutions (or rather, at the customers of these companies) and is often performed with the aid of URL spoofing, which exploits Web browser bugs in order to display incorrect URLs in the browser's location bar; or with DNS cache poisoning as a means to direct the legitimate user away from the initial target to the fake one. Once the user enters his or her password, the attack-code reports a password error, redirects the user back to the legitimate site, and poof—the user ID and password are grabbed. The easiest way to get the user to believe that the site is legitimate is to use the legitimate site (or a copy of it) as the bait.

Another form of spoofing is pharming. Pharming occurs when the spoofer sets up a redirection of a domain name from its intended IP address destination to an alternative destination in order to gain access to sensitive information (credit card numbers, account numbers, passwords, and PINs). The spoofer sets up a Web site that is almost identical to the original (PayPal is one common example) and gets people to log into the incorrect site, thereby farming—or pharming—the information that they want from the people who have been redirected. Often the redirect is accomplished by means of spam being sent to let people know that their account may have been compromised and that they will need to log in again to verify their account information. Another method of pharming can be set up to prevent the user's computer from contacting the legitimate DNS, by installing a virus on the victim's computer, by compromising the user's firewall or router, or by changing the user's hosts file so that domain names will map to an incorrect IP address. All these places can be easily obtained using information that you may have left behind during a recovery, during a test, on a set of backup tapes, or even on the server, left unguarded during a declared disaster.

Although spoofing isn't something that is necessarily more prevalent in a recovery situation, and keeping in mind that spoofing typically requires a close approximation of the "target's" Web site, the easy availability of the Web site from a recovery server may be a hole that no on considered.

Tampering The tests may uncover places where tampering may occur, either at the primary business location or at the recovery site. It is important to make note of the differences between the primary business location and the recovery site when making the distinction in tampering, and special attention needs to be paid to the places where tampering can occur at the recovery site. Although you will not likely have control over, or even access to, all the places at the recovery location where tampering may occur, anywhere you find that might fall into the category of a tampering-prone location should be investigated and noted.

There may be times or places in the recovery location where people outside of the control of the recovering organization have access to system data. They may have custody of the recovery media (allowing them to make a copy) and they may have access to the data at a low level where they could scavenge or make alterations to the data while it is within their control. It is important to research where nonorganizational people might have access to either the recovery media or the servers on which the data is being recovered and place that information in the recovery document. Although changing the data during a test would not result in much of an effect on an organization, similar alteration during an actual declared disaster could potentially affect the entire organization and its relationship with its competition and upstream and downstream partners. These situations should be noted as potential risks of the recovery site and careful note should be made during testing and actual recoveries.

Are there sufficient firewalls and DMZs in place at the recovery site? Are your cryptography keys working as anticipated? Is the security to the

network and physical security to the servers and the sever room allocated to your recovery sufficient?

If you are working on this as a small business/sole proprietorship (home based or otherwise), make sure that your LAN and/or wireless settings are adequate for the location. If necessary, get another computer (PDA or laptop) to check and make sure that the network is secure. Make note of any inconsistencies that you uncover and determine what, if anything, should be done in order to fill in the gaps.

Post recovery, these locations and situations should be investigated further and plans made to mitigate the chance that tampering can occur.

Repudiation To explain repudiation, it is important to first look at nonrepudiation. In authentication terms, nonrepudiation means that a user or a server cannot later (after they have performed an action) deny that they performed the action in question.

This speaks to the ability of an organization to ensure the security and recoverability of their cryptographic keys. Compromised keys (did you remember to scrub the servers at the recovery site and remove all traces of the organization's data and systems from the servers?) can be used by someone with less-than-honorable intentions for fraud.

FYI *Cryptographic Keys*

A cryptographic key is the piece of information that one computer uses to control its mathematical computation of the encryption algorithm to ensure that all secure messages are sent in encrypted format, (typically) with a digital signature that ensures the message originated with the sender and gets to the receiver in the same format as it was sent. This is critical to allowing a business to do business in many cases, because digital signatures are what allow a business transaction to occur digitally and remain a legally binding transaction.

Encryption with one private key requires decryption with its corresponding public key. Encryption with another key, or decryption with a different key, will result in either message text encrypted with an entirely different cryptographic key or the inability to decrypt the cipher when it is read.

Keeping your private key safe and secret is often one of the most difficult, yet most practical, issues with cryptography. Anyone who obtains the key can recover a message or set of messages, decrypt messages that are intercepted, or spoof other people by using your keys to get people to trust them rather than your business.

CONTINUED ON NEXT PAGE

This information and process is key for the purpose of nonrepudiation for an organization. An organization, as well as a person, needs to be uniquely identifiable and needs to take every measure necessary to ensure that identity theft has not or cannot occur.

Denial of Service Recall that denial of service attacks are a type of attack waged on a network in an attempt to flood the network, thereby bringing the network to its knees. Some exploit TCP/IP limitations, whereas others simply bombard a server in an attempt to render it unusable for a period of time.

Although an organization may have measures in place to limit its exposure to a DoS attack, these measures may not be entirely captured in a recovery plan and may not limit the exposure that the organization may face at the recovery site. Although uncovering this may be difficult, it is worth the effort to attempt to discover whether the measures that your organization has taken to limit DoS attacks and their effects will carry through to the recovery site. This will give you the added assurance that, should recovery to this site be necessary, your organization's existing efforts will be enough to make sure that you will be able to recover the measures to prevent these attacks as well. The results of these kinds of tests or hack attempts should be included in your recovery documents.

Threat Tree After you have identified the points where the system or systems might be vulnerable and you have determined that you may be open to a threat at one or more of those points, consider the following types of questions:

- What security mechanism do we already have in place to protect this resource?
- What security mechanism can we put into place to protect the resource?
- Are there any associated interfaces or transactions that have to be taken into account when looking at the system or the threat?

One tool that is useful in setting up the test for threats (or in analyzing them after they have been discovered) is a threat (or attack) tree. A threat tree allows you to determine the level of risk associated with each threat of attack and determine if you, as a team or as an organization as a whole, have successfully mitigated that risk.

A threat tree is a diagram, bearing an uncanny likeness to a flowchart, showing a hierarchy of threats or vulnerabilities. It shows, graphically, what might be going on in the mind of someone mounting an attack. The ultimate goal of the attack or threat is at the top of the inverted tree. Each level shows the step-by-step process that might be required to carry out the attack. Figure 9.1 is a simple threat tree for someone who might be considering stealing a bicycle.

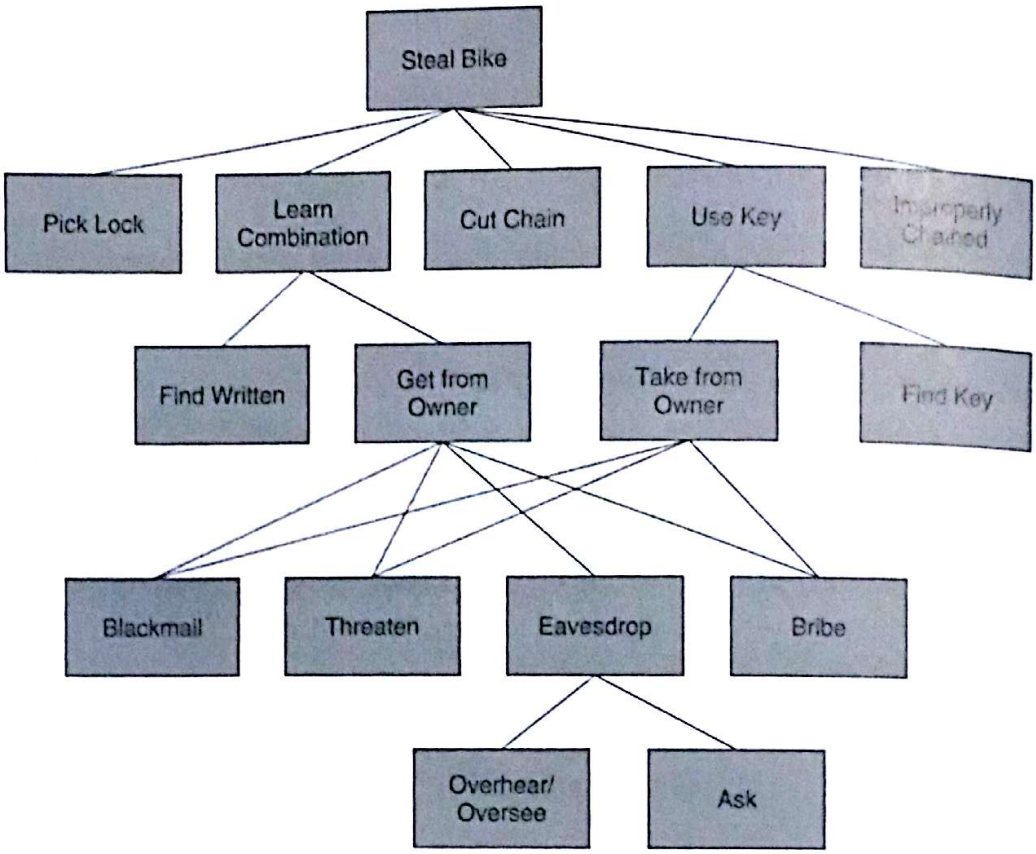

FIGURE 9.1. Threat tree for stealing a bicycle.

Outline An alternative for the graphic flow diagram for an attack tree might be to take the same kind of approach, but rather than using a flow chart outline the same details. The bicycle theft example in an outline might be as follows:

1. Steal bicycle
   - A. Pick the lock
   - B. Learn the combination
     1. Find it written down somewhere
     2. Get it from owner
        - a. by blackmail
        - b. by bullying or threatening
        - c. by eavesdropping on owner telling someone
          1. overhear or oversee
          2. just ask
        - d. bribe the owner to tell
   - C. Cut the chain (probable)

D. Use a key
   1. Take it from the owner
      a. by blackmail
      b. by bullying or threatening
      c. by bribery
   2. Find the key
E. Find the bike improperly chained

Notice that the flow chart entails less typing.

SWOT (Strengths, Weaknesses, Opportunities, Threats)

Another useful tool for disclosing not only where there are threats to the organization and the recovery process but also the strengths and weaknesses in the plan, as well as opportunities for improvement is SWOT analysis.

Typically organizations conduct SWOT analyses to determine where they stand with relation to their competitors or to the market as a whole. Marketing classes typically focus on SWOT analyses as a unit and time is spent looking at the different strengths, weaknesses, opportunities, and threats to the organization from competition. But if we look at anyone or anything attempting to break through our security measures, either deliberately or accidentally, we can apply the same kind of thought process to determining if our recovery plan and our testing scenario is sufficient.

FYI *Traditional SWOT Analysis*

Traditional SWOT analyses are tools used by organizations to audit the organization's environment. Typically, they are the first tools used by marketers to target key issues. SWOT stands for strengths, weaknesses, opportunities, and threats. Strengths and weaknesses are traditionally internal factors, and opportunities and threats are external factors.

Figure 9.2 is the graphic that typically accompanies a SWOT analysis. The further away from the center line a situation is, the further away it is from the opposite idea. The further away from a weakness a strength is, the stronger it is. The further away from an opportunity a threat is, the bigger the threat. Sometimes it is helpful to plot the location on the squares where you think each "thing" is. This will help you to see where they lie in relationship to each other and visualize how to mitigate them into the center.

>> CONTINUED ON NEXT PAGE

CONTINUED

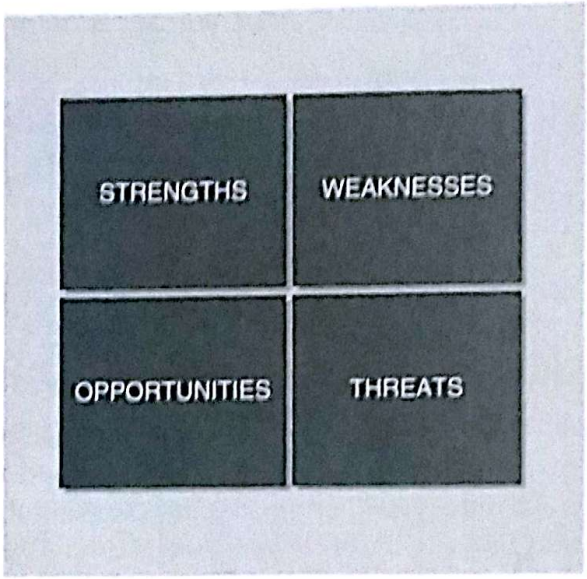

FIGURE 9.2. SWOT analysis graphic.

Strengths The S in SWOT stands for strengths. What things were done particularly well in recovery planning and recovery testing? What steps were taken to mitigate the risk to the organization from a disaster that addressed a particularly sensitive area of potential disaster? These should be addressed in this part of the analysis.

Typically, strengths point out places where an organization might have core competencies or a competitive edge over the competition. Looking at it this way, what are the core competencies of your organization's plan? What special precautions or special techniques did your team employ to make sure that you were able to recover a particularly difficult system, or what technique did you use in a challenging situation in the recovery process that allowed you to overcome the situation?

Weaknesses Where there are strengths, there are also weaknesses. What did your organization do that didn't work as well as it should have? What was missed as a potential risk? Were any of the systems unable to be recovered in the appropriate time?

If you undo all the thought processes that you employ to determine the strengths, you can better think through what the weaknesses might be. It is important that thinking through the weaknesses is done in neither a judgmental nor an accusatory manner, but only as a tool to find those situations that need to be addressed so that in the future you can either overcome the situation or put processes in place so that the weaknesses are less weak, less obvious, or eliminated and even made into strengths.

A weakness might be found to be a recovery process that is entirely too slow. Thinking through it, new ways of backing up might be used to better facilitate recovery. If you take the people involved in the process out of the equation and limit your analysis to simply the facts of the situation, it is easier to look at objectively and to turn it around.

Opportunities Often, in the process of planning, recovering, and repeating the cycle, people will have what might be termed as "Aha!" moments—times when something that no one ever thought of, at least as it relates to the given situation—can be used to make better use of resources, speed the recovery, or allow the group to leverage the economies of scale or scope.

Someone might have read something in a magazine on the plane on the way to the recovery site, or on the way back from the recovery site, they think is something the organization might be able to leverage to their recovery advantage. Perhaps there is an emerging technology that the organization might be able to make use of, or perhaps the process might work better in a different order.

These thoughts should be jotted down—on a piece of note paper, in the margin of the recovery plan, on the back of a napkin—and brought into the notes that accompany the recovery plan in the post-recovery meetings. Although some of the opportunities may not end up in the planning and recovery scenario, if they are never considered they certainly won't be. Making note of an idea that might make the recovery quicker, more efficient, or less prone to human error might allow similar thought processes to occur later when the analysis phase of the post-recovery debriefing occurs. Many times people have wonderful ideas in the heat of the moment, in the midst of a trial (like the recovery testing), that they will have forgotten when the stress is less or when they are in a different situation. These might be tragically lost opportunities to make the recovery process better if they are not noted somewhere permanently.

Opportunities are often realized simply because someone has had the temerity to think outside the box. This should not only be tolerated but also encouraged. It is one reason why many people from different parts of the organization should be brought into the process as often as possible. Although younger members of the organization are often overlooked in the planning process because they don't have as much experience with the particular systems in question, they can often be a wonderful source of new ideas and innovative thinking.

Threat Finally, we look at the threats—threats that have been considered or threats that have come to light because of the recovery testing and planning. Some threats may not become apparent until the recovery testing exercise begins and people start to think through what it is going to take, from a practical perspective, to recover the organization.

Threats might be situations that endanger the organization's security or cause hardware or software failure, which become apparent only when the recovery testing is in process. It might be an inadequacy in the recovery plan that is discovered only in the process of recovering. It may even be a deficiency in the backup strategy that ends up becoming an issue when the recovery is in process.

Solution Determination

Once you have discovered your latest list of threats, either as an ongoing security practice at your organization or as a result of your latest recovery test, you have to decide what you are going to do about them. Ideally, every threat should be addressed immediately, as it is a gap, a place where your organization may find that it is in danger of a future disaster, or at the very least an emergency situation. However, realistically, many of these uncovered threats will find themselves relegated to the back burner in the organization or discounted as being trivial or infrequent enough that they are simply not going to be dealt with.

Threats, like risks, can be weighted and categorized. Cost-benefit analysis can be done on allowing the risks to remain or on pursuing a solution to them. Many threats, particularly the ones uncovered during actual recovery attempts, will likely have to be dealt with. It may mean simply changing your backup strategy or including things in the backups that weren't there before. It may mean that additional testing will need to be done during the testing phase of the recovery.

Damage

What is the potential damage to your system from the threat? Will there be data loss in the system? Will there be compromised data in transmission? Will there be breached information? Is there potential for hardware damage or media failure?

Reproducible

Is the threat reproducible in the primary location? What is the chance that the threat will succeed in the primary location? What is the chance that the threat will succeed at the recovery location? (What is the difference in these two chances?) What are the chances that someone will stumble upon the threat or gap?

Exploitable

How much effort will have to be exercised to find the threat? How much benefit will there be if it is found and exploited? Don't think that businesses are the only ones concerned with cost-benefit analysis. Anyone trying to

intercept data or wreak havoc with your systems and their data will try only as hard as it is worth trying for the benefit that they are liable to realize from the success.

Users/Systems Affected

How many people or systems can be affected by the threat if it comes to fruition? If the threat is that someone could use your Internet connection on your wireless network, but there is no data that they can access and there is no harm they can cause other than making use of the connections, the users affected may be minimal. However, what can someone do with a pirated Internet connection? Can they use your connection to bounce to another connection, and thereby point the finger at your connection when they wreak havoc on another system? Be realistic, but understand that what you view as minimal impact to your organization might be far more extensive than anticipated.

Discoverable

How difficult would it be for the threat to be discovered by someone who really wants to find it? Backdoors in some operating systems are well known by some people, and they are more than capable of testing systems to see if the backdoor is open. This means that the threat is very discoverable. However, an internal system that is used by one or two people, not connected to the network, and has one or two default passwords in some of its systems might be nearly impossible for anyone to detect and therefore is not very discoverable.

Summary

This chapter demonstrates how continued diligence is necessary in making sure that an organization's changing needs are taken into account and how planning needs to change along with them. We have found where we might be able to leverage processes that we have in one area to make another more efficient, and we have determined that thinking outside the box is just as important in a disaster recovery perspective as in any other. We have seen how to plan and execute a disaster recovery, whether as a drill or in actual fact. We can now go into an organization and become a productive member of the planning and recovery team.

Through practice (either in an organization or in our daily lives as users of systems and data), we can enable ourselves to think more critically and more productively over time so that we can make sure that (from term papers to the systems of Fortune 500 organizations) we can recover quickly and efficiently.