Chapter 5

Identify Data Storage and Recovery Sites

Chapter Objectives

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

- Determine the best way or ways to back up your data so that it can be recovered later.
- Evaluate your offsite storage options.
- Acknowledge information as well as hardware and software as an asset.
- Determine recovery site options.
- Examine recovery site types.
- Develop recovery site selection criteria.
- Outline a recovery solution.

Introduction

Once you have determined what you need to restore, you need to have an idea of where you will be able to restore it. Depending on your particular hardware configuration, this decision may be more difficult or easier than you anticipate. This chapter will help to determine the alternatives and the wisest choice for your situation.

Data Backup

Backing up the data—how, when, and how often—is one of the first considerations that come into play in a disaster recovery plan. Many decisions need to be taken into account when you are looking at a backup strategy, more than can be covered here. However, we can look at the main points to take into consideration.

It is important when doing your disaster recovery testing that you test with a normal backup. Don't schedule a special backup just for the purpose of using it for DR drills because you believe that you can restore from it. It is important to use ordinary backups for testing.

How to Back Up Your Data

There are many ways to back up data. You can opt to invest in enough storage to maintain a backup online, on a hard disk. You can look at burning data to a DVD or CD, depending on the amount of data to be backed up and the recovery method that you end up choosing (although that choice is often made based on the type of backup rather than vice versa). Or you can follow the time-honored tradition of backing up to tape. Let's look at these different options in the light of recovery and recovery ability.

If you back up to disk, and the restoration is to the same server from which the backup was taken, recovery is simple, elegant, and quick. Disk-to-disk recovery is the quickest method because your access times do not include the time that it takes to transfer data from another medium. Looked at through the lens of disaster recovery, however, if you store your backups at the same location as the data that was backed up, and a disaster makes this site inaccessible, then your backups are inaccessible as well. Further, if there is a hardware failure or a catastrophe that destroys the hardware on which the backups are stored, then there is no way to recover.

If you back up to tape, you have a backup methodology that has been proven over time. The restore time is longer than for disk-stored recovery because the tape drive needs to access the media, and you have to unspin the tape to retrieve the data from it. Restoring an entire tape, as would likely be the case in a disaster situation, is less labor and time intensive than recovering a single deleted file. This calls into play the definition of disaster that you are using in each case. If the loss of a partial file system constitutes a disaster, but not a disaster that requires restoration in an alternative location, then there are different concepts to take into account.

Please note that the types of backup methods as well as how they are deployed are covered later in this chapter.

IN PRACTICE: Storing the Backups

Another consideration to understand with the use of tape backups is tied, to a great extent, to the location where the backups are taken and stored.

A company in Texas found itself unable to guarantee that it could restore from tape backup because the company that it contracted with to transport its tapes from its data center to the secure offsite location where the backups were stored hired new people and found that one of their carriers, rather than delivering tapes to the secure, climate-controlled location, was carrying the tapes in the trunk of his car in July and August.

A company in New Orleans chose to back up to tape and to store the tapes offsite with a company that chose to store the media in the basement. Drying the tapes out after that company experienced a flood caused by a broken water pipe was a time-intensive job, and one that no one was sure was not going to destroy the tape's ability to be read and restored from.

Another option is to make backups to CD or DVD. This media has faster seek times than tape, historically is less volatile and prone to degradation than tape, and is more portable than a SAN or a drive array. The disadvantage, however, is the amount of data that can be stored to the media. The DVDs we are referring to have much higher capacity than those that you write to from your home computer (9.4 GB compared to the 4.7 GB that you likely find with DVDs created for home use), but when backing up terabytes, petabytes, or exebytes of data on large systems it will still take a considerable number of DVDs (but then, it would take a large amount of anything to hold that much data).

DVDs are good for backing up the data for home offices or small businesses. It is important, however, to note that they are not nearly large enough to back up large amounts of data. Tape backups are still the most common method of backing up large amounts of data. Network-attached storage is also a method that is growing in popularity. Although network-attached storage doesn't necessarily meet the needs of an offsite recovery effort, it is still the backup method of choice for many large organizations.

Backups that are taken to disaster recovery testing should be the same backups that are done for the purpose of production recovery, not special backups taken for the purpose of recovery testing.

When to Back Up Your Data

When you schedule your backups (either manually or automatically) is a matter of complying with company requirements and with what you are doing and how. If you need to shut down all the systems on your server in order to take an adequate backup, then scheduling the backups for a Tuesday afternoon would

probably not be a good idea. However, if you can leave all systems up, active, and available while you are taking your backups, this may be a viable time to perform the backups. It would be advisable, however, to schedule them when the system is less active and has fewer processes running and fewer users accessing the data.

Although it is always a best practice to perform backups when servers are least active (typically after hours), in a truly global organization there is often not enough slow time for the server to allow for this. Research on server load for most servers in the organization is advised so that backups can take place during the lowest average load.

Keeping all this in mind, it is also important to remember that there are some systems that either cannot be backed up when the files are open and being accessed by the applications or that must have special processing in order to have the backups be consistent and in a manner that allows for a successful restoration of the system.

The most obvious example of this is when working with databases. Many database management systems are structured in such a way that the database has to be shut down in order for backup and subsequent recovery to be successful. Either the database has to be shut down or special settings have to be set in the database to allow for accurate backups. If these procedures are not followed, the database likely will not recover to a point where it can be opened.

Various techniques can be used to back up data that is in use, or to minimize the downtime to seconds or minutes to facilitate the backup, including snapshots or third-party utilities.

How Often to Back Up Your Data

How often you end up taking your backups ends up being a combination of math and company policy. Many companies have a records retention policy that dictates how long you must retain records and be able to recreate those records. Many governmental regulations come into play with these policies. To take into account all factors when making these determinations, you should consider the following parameters:

- The ***recovery point objective (RPO)*** is the point in time to which systems and data must be recovered after an outage as determined by the business unit. Your backup methodology must take into account the need to lose no more than an hour's worth of data, or exactly no data loss in case of disaster or the ability of the organization to recreate a week's worth of data after the systems are fully functional.
- The ***recovery time objective (RTO)*** is the period of time within which your systems, applications, or functions must be recovered (either because of contractual obligations or because of organizational requirements to withstand the loss of productivity) after an outage (e.g., one business day, 72 hours, or 2 hours). RTOs are often used in combination with service level agreements (SLAs) as the basis for the development of recovery strategies, and as a means of determining whether to implement

the recovery strategies during a disaster situation. This parameter ties in closely with SLAs and with the expectations of the organization. If you must have your business fully functional within 72 hours, then you can't rely on a backup methodology that takes 60 hours to restore because that leaves you only 12 hours to get both personnel and media to the recovery location and for troubleshooting and problem resolution.

- The ***maximum allowable downtime*** is the absolute maximum time that the system can be unavailable without ramifications to the organization, either directly or indirectly.

Where to Store Backups

This is one of the most important decisions surrounding backups that you will make. Considerations need to include the accessibility to the backups when needed; the time it takes to retrieve the backups; and the climate, geography, and topography of both the primary business and backup storage locations (whether or not these are the same location).

On Site Storing your backups on site makes for easy access to the organization. This has advantages and disadvantages. Having the backups handy means that you can restore a single file whenever someone needs it to be restored. It means that, in the case of failure, you don't have to factor in the retrieval time for the media from an alternative storage location. You can simply retrieve the files whenever necessary.

However, there are two disadvantages to this scenario. First, storing the media on site means that you have added security precautions to take into consideration, as you need to store the backups in a location that is inaccessible to most people in the organization. Retrieval of the media for any reason other than to rotate or destroy the backups at the end of the retention time needs to go through the process of getting permission from the disaster recovery team captain or whoever is responsible in your organization for such a sign-off. In this way, you can ensure that no one retrieves the backups for nefarious purposes. Second, storing the media on site means that whatever disaster may befall the organization is likely to befall the backups as well. Although this may not be true for a Denial of Service (DoS) attack, virus, malicious data tampering, or hardware failure (by far the most common "disasters" that you will have to deal with), in the event of a fire, flood, or other physically disruptive disaster you will likely be no more able to access the backups than you are to access the primary system.

Often, organizations use climate-controlled fireproof, waterproof, and even tornado-proof safes or shelter structures to store tapes on site. This lessens the risk of the tapes being damaged in a disaster and still allows for rapid access by the organization when necessary. It also allows the organization to guarantee the chain of custody and minimize the points of access to the data that is contained on those tapes. This can be of particular interest for legally sensitive data.

Offsite Offsite storage, on the other hand, is more expensive because you have to pay rental and service fees for the offsite location, as well as any services associated with that storage. There are many levels of offsite storage to consider. Depending on the size, budget, and complexity of your organization, one of these solutions should be optimal for you.

Owned offsite storage, depending on the complexity of the organization and its requirements, could be as little as storing the backups at the home of a responsible individual or at a rented storage locker. Each option brings with it advantages as well as disadvantages. Storing the backups for an organization at the home of a responsible employee has the advantage of very limited cost, but it also means that this person has to be available should a disaster occur, and also that the backups are stored in such a way that a disaster at that individual's house cannot destroy the organization's tapes.

If the organization has more than one physical site, even separate physical sites within the same town, the tapes can be stored in a vault at the secondary site. In some cases a site may have another building at the same complex, one that is an acceptable distance from the primary site or main server building, that could serve as an alternate site. The organization needs to determine the acceptable distance between server locations and the offsite location. It is important to note that this can also be tied to the RTOs and the SLAs previously mentioned.

Depending on the criticality of the data and the need for assurance of both safety and accessibility, greater distance may be required (1 mile, 5 miles, 500 yards, half a state away). Only the organization can make that kind of decision, and then only after careful consideration. These decisions need to be made on a case-by-case basis.

IN PRACTICE: It Is All a Matter of Location

Depending on the organization's location or the resources available to it, it is possible that there may be opportunities to own the offsite storage location for the backups. In many locations there are old, abandoned mines that are sufficiently climate-controlled. These include limestone, coal, iron ore, and salt mines.

Underground storage initially became attractive as secure storage of documents and people during the Cold War when people were more interested in the end results of nuclear attack. The fear of bombs has lessened, but the massive underground structures live on. They have since been transformed into locations for tape and other records storage.

(Continued on next page)

>> CONTINUED

These facilities, 200 or more feet underground, are highly secure and access is easily controlled, with only one way in and one way out, and they are virtually immune from nearly any natural disaster (fire, flood, hurricane, or civil disorder). This makes them very attractive as storage locations (either owned or contracted), and organizations choosing these locations can rest assured that their backups will be safe from nearly any eventuality.

FYI *Contracted*

One thing to remain aware of with the hiring of an offsite vendor is that the vendor must be able to assure you that it will be around when the time comes for you to declare an emergency, and that it can deliver what is promised in the timely fashion that you need it.

Don't just rely on vendors telling you that they can. See if you can get references from other people who have used them in the past. See if there is proof, one way or another, of their ability to deliver. And above all, get it in writing that is legally binding.

The contracted organization needs to be in the business of data archival and retrieval or the contract will likely not work in the end. SLAs need to be established with the vendor to ensure the organization's RTOs can be met successfully and in a timely manner.

Information as an Asset

This is all well and good, when you are looking at hardware as an asset, but it is important to remember that people and information are also assets of a company, and that companies often overlook these critical assets in DR planning. They consider, at great length, what it means to have duplicates of their data, but it is not always the hardware that fails.

Often it is software that fails, and software failures often go undetected for extended periods of time. Where hardware failure is very apparent and gets noticed by virtually the entire company, bugs in software can cause insidious corruption that may go unnoticed for days, weeks, or even longer.

IN PRACTICE: Homegrown Applications

Software often is not designed by an external company (such as IBM, Oracle, or Microsoft) but is designed and implemented internally. Sometimes, although precautions are taken and testing is done, the inevitable happens, and a set of circumstances come together that causes a calamity to the business.

For example, a very large organization had a homegrown purchasing and accounting system. This system processed purchase orders, receipts, and invoices daily. The system, with few changes, had been in place for over a decade. Receipts were processed every night, via files handed off from each division to headquarters. These files went through several COBOL programs before they were added to the database and before the costs for the receipts of purchases were attributed to the different plants.

One of the plants was in Northern Minnesota, and due to connectivity issues several days passed during which no files were handed off. When they finally were, an extremely large file was in the batch to be processed. Mainframe processing often bases the size of output files on the average file size for a given period of time. When the allocated file size grows larger than the allocated file, the program running abends (abnormally ends), and the operators restart processing after adding additional space to the output file. Ordinarily processes are put into place to account for the restarting of programs and the condition would be handled efficiently.

This time, however, the file was big enough that the program had to be restarted several times, and there were no precautions in the program to skip the already processed receipts and move ahead to where the program abended.

The run started, ran for several thousand receipts, and stopped. Operations added space to the file and restarted the program. The first several thousand receipts processed again and then several more thousand receipts processed before the program broke again. Operations added more space and restarted the program. The first several thousand receipts processed a third time, the second several thousand processed for a second time, and the batch finally ran to completion.

A sequence number that was generated whenever the program ran is the only thing that made these receipts unique.

No one noticed that the database and resulting files were incorrect for over a week, because the accounting departments at the

>> CONTINUED

CONTINUED

plants were understandably concerned when their anticipated and realized costs were radically different.

Recovery in this case couldn't be to restore the system and reprocess—too much time had passed and too much processing had been involved. New programs had to be written to determine how to back out the duplicate and triplicate receipts and not affect any other part of the system.

Not only was this a lesson in recovery and the differences in recovery, but the differences in the types of disaster from which you have to recover (in this case, internally programmed software failure) and the ways in which you can recover. Further, it was a lesson in the fact that, no matter the size of the company and the amount of effort put into disaster recovery and the effectiveness of backups, there is always something that you can learn about your system.

It is important to remember that part of the information that many organizations forget to bring to the disaster recovery site is the information that will be necessary for restoring systems and data. It is easy to remember that you need the backups of the data, but you may need information on the licenses that you have purchased for software or software-hardware combinations. This is particularly true for systems like IBM mainframes, where the license is tied almost directly to the hardware and alteration of where the operating system or other software is installed will mean that the software cannot be used. These vendors fully understand the need for disaster recovery and disaster recovery testing and will work with you to provide temporary licenses for new hardware configurations, but having your existing license key available at the DR site will mean that you can more effectively and efficiently work with your vendors to expedite the process. Time is critical in a disaster situation, and the time that you save could mean the difference between meeting and missing SLAs, and the cost can be extensive.

Metadata is just as important as licensing information. The company must determine what data is important and how important each different piece of information is. Wasting time early in the recovery process recovering information that isn't as critical as other information is a less than effective use of resources. Apply your efforts where they are going to make the most difference, especially early in the recovery process (whether it is a real disaster or a drill). Part of the planning process should include attention to which information has to be available and in what order to have the data ready for the users.

Critical data is data that has to be retained and recovered for legal reasons and for the restoration of minimum work levels. Within this definition of data, it is important to note that the critical data necessary to restore minimum work levels is the data that should be concentrated upon. You have to be able to access the data that needs to be retained for legal reasons, but that data does not necessarily have to be among the first recovered. The definition of minimum work levels is situation specific, but among this information would likely be accounting information and access to purchasing, receiving, invoice processing, and payment processing functionality.

Vital data is information that has to be retained and recovered to maintain normal business activities. This information and data represents a substantial investment by the company in time and effort, and the recreation of the data may be difficult or impossible. This data may or may not be necessary in a disaster recovery situation. This is determined situation by situation (the data in question, the criticality of the restoration, and the duration of the disaster all play a part).

Sensitive data is the data or documentation that is necessary for normal daily operations of the organization but for which there are alternative sources of the same data or data that can be easily reconstructed from other data that is readily available (critical or vital data may be sources for this data).

Noncritical data is data that can easily be reconstructed at minimal cost or that has its source in critical, vital, or sensitive data but that has less stringent security requirements.

Recovery Site Alternatives

Every organization is faced with many alternatives, not only from a backup and recovery storage perspective, but also when dealing with the determination of recovery site alternatives.

Function

Disaster recovery sites function as a scaled back office for your key functions so that your company can continue its business until either its primary site can be restored to operational functionality or until your company can come up with a viable alternative for long-term functionality. A disaster recovery site is a stopgap location so that you can keep from having to close the business while the long-term solution can be accomplished.

That means that when you are looking at the different options and the different locations, you need to keep your business and its functions and needs in mind. Make sure that what you are looking at will meet the needs of the company for at least a couple of months. If you are back up and running as usual in a shorter time, that is wonderful, but you should look at everything with realistic expectations.

Hot Backup Sites A hot backup site is a location, other than your typical business site, where a duplicate of your data is stored, ready for access in a moment's notice. Real-time data transfer occurs between the primary and the

hot backup site, which means that the site has to be contracted continuously or that the company has to own both the primary and the hot backup site.

The advantage of a hot backup site is that the company can guarantee that there will be at most a certain period of data loss (for example five minutes) that can account for the lag between the two sites. This can be a core competency or a selling point for an organization to its clients.

The disadvantage, of course, is cost. Not only do you have the cost of either contracting or acquiring the backup location, but also the cost of the network traffic that is required to keep the two environments continuously in sync.

Warm Backup Sites A warm backup site is a facility that is already stocked with all the hardware that it takes to create a reasonable facsimile of what you have in your primary data center.

In order to restore the organization's service, the latest backups from your offsite storage facility must be retrieved and then delivered to the backup site. Once this is accomplished a bare metal restoration of the underlying operating system and network must be completed before recovery work can be done.

The advantage to the warm backup site is that it can be gotten to and a restoration accomplished in a reasonable amount of time. The disadvantage is that there is still a continued cost associated with the warm backup site because you have to make sure that you maintain a contract with the facility to keep hardware up-to-date with that which is found in the organization's data center. This is the only way to ensure that the backups can be read when the time comes and the only way to be sure that the restorations can be accomplished successfully.

The warm backup site is the compromise between the hot backup site and the cold backup site.

Cold Backup Sites A cold backup site is a building space that can be leased or located (if you are very lucky) when a disaster occurs. Everything that will be required to restore your organization to productivity will have to be procured and delivered to this site when the disaster has been declared. Only after the physical resources have been acquired can the process of actual recovery begin.

The advantage of a cold backup site is that the cost is minimal. Even if the organization chooses to contract the floor space long term, the cost is far less than that associated with the others.

The disadvantage of a cold backup site is that the organization can find itself unable to recover operations for an extended period of time, particularly if the organization relies on resources that are difficult to obtain or that require extended lead times to be configured. If an organization chooses to rely on a cold backup site, and it is not reliant on commodity hardware, that organization may fail to recover in a timely enough fashion to remain a profitable business and may fail.

Reciprocal Agreements Reciprocal agreements are often a viable and useful alternative to a warm recovery site, but you have to make sure that the

company you are working for and the company that you are anticipating entering into such an agreement with are not only aware of the advantages and disadvantages of such an agreement, but that the company that you work for is a good fit for the agreement long term. A company that is not a good fit long term may still use a reciprocal agreement as a short-term solution to the disaster recovery scenario, but all parties have to acknowledge that the solution is, indeed, short term.

A reciprocal agreement is the agreement of two (or more, but this is less frequent) companies to join together and act as partners in disaster recovery. One kind of business that can make use of this kind of agreement particularly well is the home office or startup business. These businesses often have shoestring budgets and can little afford to be without access to their information for extended periods of time, and can afford even less the outlay of money for the purpose of making sure that there is a place where they can recover their information should a disaster or unplanned event prevent them from using their own resources. Two or more of these companies can enter into a simple agreement so that they can share the cost of resources over the duration of the agreement.

For example, Larry is a freelance author. He is working nearly full time at writing while also holding down another full-time job to pay the bills and provide insurance and support to his family until he feels able to get into writing as a single full-time profession. Adam is a tax accountant who is just starting to venture out on his own. He, too, is working at another job until he can build his client list to such a degree that he can see his way clear to make his own business his only profession. Amandya runs a home crafting business making custom scrapbooks for customers' weddings, their children, or other significant life events. This is her only official job. Finally, Joy has an established, though limited, enterprise that provides housekeeping services to a given set of clients. All these people have their own set of information that is relevant to their particular businesses and their own needs from a hardware and software perspective. However, many of the needs overlap or are exceeded by the needs of another (one application has greater hardware requirements that others may be able to take advantage of even if the need is not present) and can therefore be easily met. Some, such as the tax accountant, have specialized software requirements that the others don't have. These four organizations can enter into a reciprocal agreement that can include the cooperative effort to purchase a backup system that is stored in a central location (another person's home, or in a storage locker somewhere) that can be used by one or more of the organizations (at one of the four offices in question, or at another agreed upon location) for the purpose of disaster recovery. For this to work well, the parties involved cannot reside in very close proximity to each other (for example, the same street or the same housing area, and likely not even the same town) so that there is a good chance of recoverability should a hurricane affect Amandya's location or a forest fire affect Larry's access to his home office.

Such an agreement is typically made up of two parts, agreed upon by all parties. The first is a letter of understanding that lays out, ahead of time, any costs associated with the agreement and a general agreement of what responses

each company expects to offer and receive during an outage. This letter of understanding should also provide information on proposed recovery sites for each party in the event of a disaster. For example, should Joy need to recover on the backup system, she would have a certain amount of time to either inform or request from any other party (Larry, Adam, or Amandya) that she is in need of temporary recovery facilities, and office space could be made available to her at one of the alternative locations. The second part of the agreement is a completed configuration questionnaire that covers current and anticipated hardware and software requirements, and an anticipated schedule for updating the software and hardware and making the associated changes to ensure that all parties' needs will be met and that all parties will be able to use the systems in question in the event of a disaster.

At some point, the parties involved may find this to not be an adequate solution or they may find that their needs and the needs of the others begin to diverge to an extent that the agreement is no longer anticipated to be beneficial. In this case, it is important that the agreement be equitably dissolved and the parties that choose not to remain in the agreement find other recovery solutions.

Another interesting situation where reciprocal agreements are profitable is when key personnel in the organization have their own resources at home from which they can perform many of their ordinary work duties temporarily. The use of these resources as auxiliary disaster recovery resources can be contracted as a reciprocal agreement with the reimbursement for the personal computer and supplies needed to be part of the disaster team. The benefit to the employee would be financial; the benefit to the company would be the readily available resources should they prove to be in short supply at some point.

Two Data Centers This alternative can be seen as a logical extension of either hot or warm backup solutions; however, it can be defined differently when the organization in question isn't sufficiently large to contract either for a hot or warm backup site or to own a hot backup site in another location. If this is the case some organizations, typically smaller ones, can equip a limited second data center.

For smaller organizations second data centers can mean (if the company is small enough) a couple of laptops equipped with sufficient hardware resources and the proper software necessary to conduct business. It can mean a duplicate set of hardware (if a couple of laptops are not sufficient for the amount of resources or if laptops are not robust enough to provide sufficient power for the organization's need) set up with software stored in an offsite location, such as a storage locker or a rented space in a warehouse.

This solution is limited in cost, but can be sufficient. If space is required temporarily to recover the business, office space is easily contracted on a limited time basis either from official locations or, in a pinch, from a hotel in a safer location.

Consortium Arrangement If we consider a consortium to be a group of organizations that have formed some kind of association or formal combination with the purpose of engaging in a mutually beneficial venture, then we can extend that to a disaster recovery consortium arrangement and see that the venture would be to provide disaster recovery facilities to the other members of the consortium. This isn't totally unlike a reciprocal agreement, but as the group is often formally joined together, the venture ends up being more formalized and longer lasting than the reciprocal agreement.

Since it is unlikely that any of the members of the consortium are in the group for the sole purpose of providing the other members with backup facilities, it is more likely that each member would provide a given amount of space on existing hardware to other members for the purpose of all or part of their temporary recovery efforts.

Because few organizations, especially in today's business environment, can afford to dedicate significant chunks of their resources to remaining available for use in a disaster situation for one of the other members of the consortium, this type of agreement is not typically used.

Vendor-Supplied Equipment Agreements are often made with vendors to be able to provide emergency equipment to organizations in the event of a declared disaster. Space and the utilities to make a functional representation of the organization are often easier to come by in alternative locations, and often are contracted for extended periods of time, but the hardware on which the systems run is contracted only to be made available in the event of a disaster.

This is much easier to do if you are requiring the use of commodity type equipment rather than custom equipment. For example, PCs, even several dozen or several hundred, are far easier to come by than two or three PDP-11s. With the increasing popularity of Linux as a reliable operating system for industry, this commodity kind of hardware is becoming more prevalent, but the reliance on mainframes has not diminished and it may be the fact that, depending on the extent and severity of the disaster, some equipment may not be as readily available, even to fulfill vender contracts, as it could be.

Caution should be taken when entering into this type of agreement to make sure that the vendors in question will be able to fulfill a need such as this in a timely enough fashion for the organization.

Combinations In disaster recovery, as in many areas of business, exactly what happens and how is not always cut and dried, not always black and white. In a lot of cases, one solution does not fit every aspect of a company's disaster recovery needs. In these cases combinations of solutions are created to custom fit the needs of the organization.

IN PRACTICE: Custom Solutions for Solutions Reference

Solutions Reference is a research company that assists companies with marketing research and with research and development based on that marketing research. It has determined that, in different situations, it needs to employ different recovery methodologies.

In the event of a natural disaster that affects the locations where it has its primary computers, the company will take its daily backups to an offsite recovery location in a different state so it can recover while its primary location is being brought back online.

The company has determined, however, that (because it is in a hurricane prone location) it can opt to head off the effects of a disaster by temporarily relocating the servers for a period of time to the servers in the winter home of one of the company managers in California. The manager simply needs to fly himself (and potentially his family so he has that piece of mind) to California, attach his system via the extended network to the network at work, and make a copy of the system. The switch over can occur at that time, meaning that there is literally no downtime for the servers and the clients are unaware of any potential for disruption.

Written Agreements

Written agreements for the specific recovery alternatives selected should be prepared, and should include the following special considerations.

Contract Duration The duration of the contract needs to be set and agreed upon by all parties. Close attention should be paid to when the end dates are approaching so that renegotiation of the agreement can occur and the parties in question are not left stranded without an agreement and a disaster pending.

This is one reason why a single owner of the disaster recovery plan and all associated materials is usually critical to an organization. One person (with a backup, of course) owning all disaster recovery material means that that person is responsible for keeping up with agreements, licenses, and contracts and that the organization can rest assured that it will not be left without an agreement should it need to rely upon it.

It is particularly convenient if agreements and contracts can be written so that they all expire at nearly the same time, so that all negotiations can occur around the same time. This will make remembering easier for the responsible parties.

Termination Conditions All contracts and agreements should have termination conditions that allow either party to terminate the agreement, although

there are usually financial repercussions for whoever decides to terminate before the expiration date of the contract.

Conditions for termination include the insolubility of one of the parties or the occurrence of too many declared disasters without justified cause in too short a period of time. Although some disasters are unavoidable, such as floods and tornadoes, others are often preventable in part or completely by taking reasonable precautions.

One example is the case of Crystal Dogs. The Crystal Dogs company was a dot-com startup that sold crystal dogs online. It was a home business retailing crystal dog figurines and Christmas ornaments, an Internet-only business whose founders ran the Web site and all associated systems. These owners contracted with another service provider (DR_Cheap.com) to be their backup for a nominal fee so that, in the event of an emergency, the business could be up and running again without having to resort to spending a lot of money on a backup site.

Crystal Dogs, however, failed to take the minimal necessary precautions to secure its network and install firewalls and filters on its systems, and it fell victim several times within a 6-month period to DoS attacks and viruses and had to rely on the services of DR_Cheap.com an inordinate amount of time.

DR_Cheap.com felt that they were losing money on the proposition (kind of like an insurance company when you have an unusually large amount of accidents in a short amount of time) and dropped the contract with Crystal Dogs because it felt that Crystal Dogs wasn't taking the necessary precautions to prevent declarable disasters.

Testing Any agreement should have, as a part of the services provided, the facility to allow your organization the ability to test your disaster recovery plan. It is important to note that, unless you have fully tested the plan and are confident in your ability to recover from a disaster with the given plan, you really can't say that you have a disaster recovery plan. You can say that one is in process, but without knowing that you can rely on the plan in case of a declared disaster, you can't say that you can actually recover.

Costs Another critical feature of the written agreement needs to be the provision for all costs associated with the agreement. The contract needs to detail anything that isn't covered by the agreement, and you need to make sure that it includes anything that you think ought to be covered. Usually hardware is covered, but are there any software costs that you think ought to be covered that aren't apparent in the agreement?

Although this may not seem to be overly important for many software products (such as Microsoft Office), there are times when not having the correct license key can mean the difference between being able to recover during a test or declared disaster and not being able to.

For example, there are keys associated with VMS and ZOS operating systems without which you cannot install the operating system. These keys are costly

and often associate a given version of the operating system with a given configuration of hardware. Deviation at any level from the original configuration can invalidate the keys. Special keys may need to be acquired in the event of a test or a declared emergency, and it is best to know upfront if you are to be responsible for acquiring these or if there will be additional costs associated with the acquisition of them should the contracted company be required to acquire them.

Special Security Procedures In many cases, the agreement will have a clause or clauses concerning security measures that are required or being provided. It is important to make note of these measures and procedures as they may or may not be sufficient in your organization's given situation, and special provisions may need to be made to ensure the safety of all parties and the security of property and information in the event of a test or of a declared disaster.

Notification of Systems Changes There should be a process set forth in the agreement documents providing a process or procedure by which you can change the configurations of your systems covered by the agreement. There will likely be charges associated with these changes; however, as a means to protect your organization's interests in being able to upgrade and stay current on systems, it is worth the extra money to have written into the agreement the clauses that will allow you to make such changes.

Hours of Operation What happens if you have to declare a disaster at 3 A.M. on a Sunday morning? Will you be able to reach anyone at the DR company to let you in and help you get set up? Will you be able to find someone to call? Will you even be able to get into the site before 9 A.M. on Monday morning? What will it do to your company if you are unable to start your disaster recovery efforts until 30 hours after the disaster happens? Will you be able to access your backup tapes in that time frame, or will you have to wait until Monday morning for that, too? The hours of availability for all agreements should be spelled out so that you can be sure you have access to the resources required when they are needed, not necessarily just when it is convenient.

Specific Hardware and Other Equipment Required for Processing Does your organization have special hardware needs? Are you running on a PDP-11 (a DEC computer from the early 1970s and still in use in many companies today) and you need to have that available to meet your disaster recovery efforts? If you have such a need, it should be spelled out in detail in any agreement. This extends to any odd tape devices required for recovery, printers that are programmed for in your application (do you only print your reports on certain-sized paper, do you need a special printer to handle your check or invoice printing requirements, or is your application set up to send commands to a DEC Laser printer or an HP printer and the company with whom you have a contract has standardized on something else?), or any other quirky hardware that may be difficult to obtain on short notice.

Personnel Requirements Many companies will contract to provide hardware and even software to another company in the event of a declared disaster or for the purposes of testing a disaster recovery plan, but what happens when you need people to do the job? First, is your disaster recovery plan sufficient that it would allow someone walking in off the street to do the job for you? Then, what are the security requirements for the people performing the recovery if they are not employed by your organization? Finally, where are you going to get these people, how much are you going to have to pay for them if the need arises (the Year 2000 "scare" taught many people, COBOL programmers in particular, the real facts about supply and demand, and looking for people who can fill roles with no warning in a disaster may mean paying an extreme premium for the service), and how much notice is going to be needed to get qualified people to perform the recovery or the tests? These questions should be spelled out in any agreements that you have that could be impacted by the answers to the questions.

Circumstances Constituting an Emergency Arguably one of the most important parts of the agreement is the definition of situations that may arise that constitute an emergency. One company's total fiasco might only be a bump in the road to success for another. Planning far enough in advance for what your company might see as a sufficient disaster or emergency not only allows you to plan but also those organizations with whom you sign agreements to plan for such a contingency.

It is important to outline what constitutes an emergency. Although you don't actually have to detail everything that might happen, you should at least make sure that you find a way to quantitatively detail what the impact to your organization needs to be in order for you to declare a disaster and therefore make use of the agreement.

Process to Negotiate Extension of Service Because any agreement is limited in duration, the process for renegotiation of services needs to be outlined in the agreement or in auxiliary documents to the agreement. This is important so that both sides of the negotiation know what to anticipate and when to anticipate it.

Nonmainframe- or Nonserver-related Resource Requirements Although many people, and indeed many companies, think of mainframes and midrange computers or Linux and Windows servers when they think about disaster recovery (some even take the printers into account), many don't think about the smaller and less glamorous resources such as desktop computers, laptop computers, phone systems, and other devices that allow people to actually do their jobs.

Because many disaster recovery service providers don't provide for the computers that the users will need in order to get back to work so that the company can continue with business as usual, those facilities need to be planned for as well.

For example, IBM doesn't necessarily have the facilities at their Boulder location to house a hundred or more people sitting at desktop computers without special arrangements being made to make the facilities available. However,

it likely would be possible to find, in the same general location, office space or warehouse space that could be made ready fairly quickly so that users could be brought to the area of the disaster recovery site for the duration of the recovery efforts and until the main office site could be made safe for business again.

Priorities Recall that in Chapter 1 we looked at breaking necessary functions into groups of prioritized jobs for the purpose of making the recovery and the recovery planning more efficient and organized. These priorities can be translated into the agreement. The resources necessary for a tier 3 recovery need not be available at the same time as those that are necessary for a tier 1 recovery. In the best-case scenario, those resources may never be required because the organization's main office may be available by that time.

If the resources required can be put off, then cost can be spread out over the duration of the recovery effort and concentration can be paid to the job at hand rather than worrying about the extra hardware that may be around the area.

If there is hardware around, someone in IT will often try to recover everything available rather than concentrating on the job at hand. If there are spare cycles, a technologist tends to try to apply those cycles to any job at hand. It is better, in this case, to pay close attention to the job at hand.

Other Contractual Issues There are likely other issues that need to be put in writing. Legal assistance should be sought with this as with any legal agreement. If you are with a large organization, there will likely be legal council already associated with the organization. If you are with a smaller company or you own the company, you may need to find legal assistance for the purpose of these agreements.

Alternative Site Selection Criteria

Whenever you are selecting possible sites for your disaster recovery site, you will be faced with many alternatives. It would be helpful to have a list of prioritized features that you are looking for in a recovery site. This section provides you with some of the criteria that are commonly used in the determination of provider and site selection criteria.

Number of Sites Available

Although you will likely want to recover to the closest location available to limit the cost of travel and the time needed to get to the recovery site, a single disaster may affect more than one location, and different disasters may happen in different locations at the same time. The availability of more than one location can mean that you have choices in the event of a declared disaster. You may have to agree on a primary recovery site so that the provider can plan ahead for the eventuality of your needing to recover at this location.

Many companies can't offer a wide selection of recovery sites, but they will be able to give you at least two or three different locations to which you can relocate your functions.

Distance from Site

Keeping in mind what your company will need in order to recover and get back up and running in a timely manner, and the time frame that you have committed to be able to do just that, you will need to determine the best places to find a recovery site. If you are in Minnesota and you have 72 hours to have your recovery done, it isn't likely that you will want to use a recovery site in Australia. Colorado may be a viable alternative, or even Atlanta.

Another reason to look at distance from the site is that if the site is too close to your business, you may find the recovery site affected by the same emergency that is affecting your business's primary location. If a hurricane were to affect Tampa, choosing a disaster recovery location in Tallahassee may not be an optimal solution. Houston or even New Orleans may not even be good solutions. Pittsburgh, however, may be a solution for you or New York or Delaware.

Balancing the need to stay far enough away to allow you to escape the disaster and being close enough to make getting to the site in a timely manner practical can take on the air of an art rather than a science.

For example, if you are recovering a business located in Amarillo, Texas, or Hibbing, Minnesota, you may need to take into account that there are a limited number of locations to which you can fly directly, even though there are international airports available. You may have to plan in time to get from your location to another location, stopping at interim airports with an indeterminate layover at any of them.

In narrowing your choices, look at current flight schedules and notice not the shortest time from point A to point B but the longest time it could take you to get there given the schedules. Then double that time to account for possible delays or cancellations. It is better to have a realistic idea than to suggest the best possible picture to the team and to upper management and be proven wrong at the expense of the company. The tighter the timeline required and the tighter the budget for the recovery effort, the more important it is to err on the side of pessimism.

Facilities

Look at the complete picture of the recovery site and the surrounding location. How hard is it going to be to get what you need while you are there? Don't just consider what will be needed immediately for the recovery proper. What will be needed for the period of time after the recovery for employees to get back to work and be productive? Computers are a given, but what about notebooks, pens, and envelopes? And people can't work 24 hours a day. What is the local environment like? Let's look at some of these considerations.

Office Supplies One of the most overlooked things in a disaster recovery plan and in the selection of the disaster recovery site location selection is the ability to acquire office supplies. Although in a perfect world we could exist in a paperless society, in reality, this really isn't much of an option. People still prefer to take notes on paper with a pen. Even a stylus and a PDA don't usually offer people much compensation for being able to write with a paper and pencil.

Printer paper and toner may or may not be included in the cost of the recovery site. You should know upfront what is and isn't provided and you should take into account where you can purchase, locally or quickly, the products that you will need.

Meals You have people; they have to eat. If you are looking at being at this location for an extended period of time (a month or more), having everyone subsist on vending machine food isn't practical or healthy. Paying attention to the restaurants and supermarkets in the proposed disaster recovery site area means that you can prepare for the fact that you will have many people who will need to be fed and in turn can prepare those going to the recovery site for what they will find there.

One company I know of created a disaster recovery journal. It is a three-ring binder with menus from local restaurants, maps to the restaurants from the DR site, and maps to local shopping establishments (malls, Wal-Mart, grocery stores) so that anyone who ends up at the disaster recovery site (for testing or for a declared disaster) can go to a central location and find information on what they need to know. Every time there is a DR test, one person on the testing team is designated to go to all the local establishments and get a new menu for the journal. This may not be a part of the true disaster recovery, but it is something that will make those involved more at ease.

Living Quarters Look at the proximity of hotels, motels, and short-term housing to the disaster recovery site. A site is less practical if it has only long-term housing (a one or more year lease) close to it.

Again, making an up-to-date list every time there is a DR plan or having someone maintain a list of these local residences will allow those traveling to have a starting point both for the testing and for a declared disaster.

Postal Services If your organization is going to have to send and receive mail (bills, invoices, payments), having ready access to postal services will be a must. If you are going to be at the backup location for more than a few days, you will likely find yourself having to rely on the postal service in the recovery location. Knowing where it is and how to access it and creating a working relationship with the people there will go a long way toward easing pains should disaster declaration become necessary. Knowing where to access and how to locate these facilities should be part of the criteria what you base your decision for the choice of location.

Recreational Facilities Okay, so maybe it isn't one of the things that people think of first when looking at disaster recovery sites, but when you are dealing with people who are a long way from home, under stress, and probably working long hours, in order to keep tempers from flaring, and so that you can make the most productive use of the hours that you do have, having recreational outlets for people will go a long way toward maintaining morale.

Whether it is a gym, park, golf course, or other outlet, knowing that there are places central to where the recovery site is located to give people a place to work off stress should be on your list of things to look for when looking at recovery sites. The company is going to have the company as its top priority. This is a given. But in order for the people involved to be as productive as possible, providing them with this outlet will do more for the bottom line than anyone could anticipate.

Cost

In the long run, it frequently comes down to cost. This may not be the only deciding factor, but it will likely be one of the primary factors leading to the decision.

Site Cost The site, regardless of where it is or what it provides, is going to require a significant outlay of money. But the organization needs to look at this more as an insurance policy than as a sunk cost. An insurance policy is something that you buy and pay for on a regular basis, something that you hope you never have to use but that you know is there in case you do. The same holds true for the disaster recovery site. Although you may make use of it on occasion for testing (kind of like health insurance, you use it for physicals), you hope you never need it for anything really serious.

Travel Cost Travel cost may be even more of a consideration than the monthly cost of the contracted recovery site. Depending on where the recovery site is and how air travel works in that area, you may be looking at higher-than-usual expenses for the cost of tickets. But that is only part of the travel cost that has to be considered.

The cost in man-hours should be considered as well. A quick look on Priceline (**http://airlines.priceline.com/airlines/flights/flights-to.html**) might allow you to see the following information besides the price of the ticket. With this information you can infer more information that may impact the decision of the company as to which recovery site to take.

For example, if you are working for a company in Clarendon, Texas, the closest airport is in Amarillo. If the recovery site is in Boulder, Colorado, it could take you an hour to get to the airport, up to five hours, if luck is on your side, to get to Denver (the closest airport to Boulder), and then an hour or more by car to get to the hotel or the recovery site. You could drive nearly as quickly.

And it gets worse. If you work for a company in Lubbock, Texas, it could take you eight hours to fly to Denver, and many carriers route you through Dallas and Chicago. Add the drive time to Boulder, and the time it might take you to get to the

airport in Lubbock, plus the wait time for the flight because you never know how close to the flight time the disaster will be declared, and you could be chewing up a significant number of man-hours, not to mention recovery hours, in the process.

These are all hours that you can't spend recovering your data or your systems. But they are hours that count against those hours in your SLAs.

FYI *Air Travel*

Keep in mind that, in the event of a declared disaster, you won't be able to rely on getting the best prices from somewhere like Orbitz or Priceline. You will likely have to walk up to the ticket counter at the local airport and pay whatever the asking price is for the next flight out. And you may not be lucky enough for that flight to be leaving right when you want it to. You may find that you will have to wait extended periods of time for the flight that you need and pay a premium for that wait. All these costs are costs associated with the choice of recovery site.

Cost of Temporary Living No employee, no matter how good-hearted, will foot the bill for the cost of temporary living expenses incurred while they are either at a recovery test drill or at an actual recovery. The company will have to compensate the employee for the temporary living costs. The more people involved, the more that expense can add up. This is yet another expense to be taken into account. The cost of living in Clarendon, Texas, is likely to be somewhat less than the cost of living in Boulder, Colorado, and the cost associated with living in Boulder during peak tourist season for extended, though temporary, periods of time can get to be excessive.

This is not usually a cost that is considered in association with the location of the disaster recovery site, but it ought to be taken into consideration.

Contract

Naturally, the contract should be considered when looking at the recovery site. Terms of the contract, what is included in each contract that the company is evaluating, the duration of the contract, and any additional costs for goods or services that may not be included in any given contract should be considered.

Designing Recovery Solutions

Without the ability to recover, there is no purpose for any backup other than to waste media. Going hand in hand with whatever backup solution is in place, or that you are putting into place, is the knowledge that at some point you will likely have to recover using these backups.

We have to establish a recovery site, select backup and recovery strategies, identify the tools necessary to meet our storage needs for the backups we have taken, and identify those places where creative solutions need to come into play.

Establishing a Disaster Recovery Site

One of the first decisions that has to be made in the disaster recovery planning process is how you intend to structure your backups and your planned recovery. Will there be automatic failover and no intervention will be necessary? Or will you have to retrieve tapes from offsite storage and have them transported to an alternative recovery site? Each alternative comes with its own advantages and disadvantages. Each also comes with trade-offs and costs. Some costs are direct, as in the physical outlay of capital, and some are indirect, as in the cost of bodies and time to recover the organization.

Choosing a Site: Hot, Warm, or Cold Standby The choice of sites often comes down to a matter of money and practicality. Many companies would rather have a hot backup site, where they can fail their entire business over within a matter of minutes to hours. This kind of assurance, however, costs money—typically a lot of money. It means providing an entire office complex in another location (almost always another city, and usually another state). If the company is a manufacturing company, this could also mean the contracting of another facility or subcontracting work out in case the primary facility becomes unavailable.

This choice may be influenced by clients, stakeholders, and shareholders, but ultimately it is a matter of measuring the return on investment and determining how long your organization can withstand the risk of being without its ability to process work.

There are cases where a hot backup is necessary. Recall that we undertook a risk analysis in Chapter 3. We can begin to think in terms of identification of that risk as it pertains to the identification of recovery solutions. The acceptance, or the non-acceptance, of risk will drive many of the decisions that the organization makes surrounding the backup and recovery solution decisions. Recall that risk is the chance that an event will occur. This chance of an event occurring, coupled with the loss of revenue that will occur if that event occurs and the cost that is associated with it, are all factors that have to play into the organization's determination of whether the cost of the recovery solution is justified.

If you determine that there is little chance that a disaster or even an emergency event will occur, then you have minimized the risk from the beginning, and the choices that you make concerning recovery solutions is very much simplified. If, however, the organization determines that there is a significant enough risk surrounding the organization, there will be additional amounts of investment that the organization will have to be willing to make.

IN PRACTICE: Small Business

Risk, and how an organization chooses to deal with mitigating that risk, is relative. Big businesses tend to be able to throw more money at a solution than smaller businesses, but the size and complexity of solutions is relative. They may not be throwing any larger a percentage of their money at the solution than a small business or a home-based business might be in similar circumstances.

Small businesses are sometimes less adaptable to disaster than large companies. Someone who has turned a room in their house into a small-scale manufacturing line may not be any more or less able to withstand the occurrence of a disaster than Ford, Dodge, or Westinghouse. If all the capital and thought is tied up in creating the product, and little is tied up in learning what they would have to do if a disaster levels their house, they are in little better position than a large organization in a similar situation.

However, an accounting company that is based out of the home office of the accountant may be totally recoverable with minimal downtime if that company invests in the time and material necessary to make a periodic backup of the software and records on its system and stores those backups in a waterproof, fireproof file cabinet in a closet and makes two duplicate copies (DVDs are even getting cheaper today) and sends one to a relative in Pennsylvania and puts the other in a safety deposit box in a local bank.

Simply replacing damaged or destroyed hardware from a local commodity hardware store (electronics store, office supply store, or even from the mall) and recovering the backups would be all that it would take to recover the business in a storefront three miles from the primary location or in an apartment or hotel room nearby.

It is up to every individual organization to assess the risk and determine what it will mean to that organization and how best to recover should the worst happen.

Build vs. Rent or Share The decision over whether to build a recovery site (whether dedicated to nothing but being a recovery site or as a line of business of its own) or to rent or share a recovery site with other companies is often a difficult one. The level of trust that has to be involved between two organizations in order to share a recovery site implies a special kind of relationship. Although this kind of relationship may exist fairly easily between two home-based businesses where there may already exist a trust relationship, the same level of trust may not exist between two larger organizations. Not only is there a level of trust that has to exist, but there has to be a small chance that both organizations will need to use the recovery site at the same time and that both are going to be able to uphold their end of the agreement over time. On the other hand, it is typically more affordable to establish a share relationship than to shoulder the entire cost of the recovery site. In the end, the decision to rent or share comes down to total cost.

Selecting Backup and Restoration Strategies

The choice of backup strategies in order to make the optimal use of both time for backup and time and effort for recovery is often a decision that will be based as much on guess work and art as it is on facts and figures. Many times the backup strategy that is chosen by a company is simply the favorite of one of the administrators or is just something that everyone seems comfortable with. These, however, are not valid and sound justifications for a backup strategy.

Full backups are complete backups of a system. They are a snapshot of how a system looked at a given point in time. Full backups (often referred to as cold backups) are invaluable to an organization for many reasons, and should always be a part of any backup and recovery solution. They may not, however, be the all-inclusive solution to a company's every need. In order to take a full backup, all the applications and databases on a system need to be shut down. This means that time needs to be taken to complete the shutdown and the restart after the backup is taken, and a backup taken of the full system. The disadvantages of full backups are that they take time and that you can only recover to that specific point in time, if that is the only backup on which the company is relying.

Incremental backups are those backups that gather together, in a perfect scenario, only those changes that have occurred since the last full or incremental backup. This kind of backup is not always possible without the assistance of third-party backup tools to assist with the backing up of the data and surrounding application. Incremental backups provide the fastest backup scenario, because you are only backing up the data that has changed in the period of time since the last incremental or full backup. What's more, because you are backing up the minimal amount of data that is practical, the storage space for these backups is the smallest. The trade-off here is that your recovery scenario can be longer because you have to incrementally restore each of the files from the original full backup to the point in time to where you need to recover. If you don't take full backups very often, this can be extensive.

Differential backups are those backups that contain all changes since the last full backup. The advantage of a differential backup is that you only have to restore the cold backup and one additional backup in order to restore the system to any given point in time. The disadvantage is that you have to back up and store the backups of redundant data. For example, if you take full backups every Saturday night, and differential backups on Sunday, Monday, Tuesday, Wednesday, Thursday, and Friday, by Friday you have duplicated the data that you backed up on Sunday five times.

Mirror backups are those backups that directly copy all selected files, directories, mount points, and file systems from one set of disks to another. A mirror backup is conceptually identical to a full backup except that the data in question cannot be compressed in any manner and cannot be password protected. Mirror backups are often made when using RAID by splitting the mirrored device and maintaining one of the mirrors (or backing up one of the mirrors) as the backup. The space required for a mirror backup is equal to that of the data that it is backing up, but typically (as long as the backup stays on disk) the recovery from failure is faster. This solution, however, is not typically practical in a disaster situation because a disk array is difficult to store and transport and a disaster at the primary site will likely wipe out the backup as well. Rather, a mirror backup is normally used as the source for an incremental, differential, or full backup so that the backups can be carried out without impact to the source system. The requirement for a duplicate set of hardware (or triplicate, depending on the RAID configuration and redundancy requirements of the organization) remains, but the recoverability of the system is greater when the mirror is used as a backup source rather than the backup itself.

Most organizations will come up with their own combination of these backup strategies. They may delineate their backup strategy based on the amount of activity on a system or on the features available to a given system by the underlying architecture. It may be easy to back up the entire system weekly with a full backup and then forget about most of the executable programs because they don't change much. The data in a transactional database (a database that has users inserting data into and altering data already resident in the database) will likely change minute to minute. That data will need to be backed up frequently, either incrementally or differentially (the choice of how has to be based on your organization's recovery requirements), in order to be sure that you can recover the data as quickly as possible to any given point in time. The data in a data warehouse, because it is far less volatile, may change only once or twice a week, depending on your load strategy, thus a single differential or incremental backup after those loads would be sufficient and similar to each other in recoverability and storage.

Hot backup is another alternative that is becoming more popular with larger organizations. With a hot backup, the entire decision (backup site location selection, backup and recovery strategy, vendor, etc.) is a package deal because the recovery location is in constant communication with the primary systems. This means that, in the event of an emergency or disaster, connectivity to the systems automatically fails over to the backup site and there is effectively no downtime noticed. This is an optimal solution for many organizations that have multiple data centers or multiple locations that can be used as data centers.

Although any backup solution has to effectively fit within the window available to you for the backup, the primary consideration in these different options needs to be the recovery needs of the organization. The most elegant backup strategy is worthless if you cannot effectively recover using the backup

in a timely manner, and in a manner that meets all the needs of the organization. It is not practical to believe that, if it takes 10 hours to take a full backup of a system and each incremental backup takes 5 hours, and you need to restore a full backup and five incremental backups, that your recovery window will be any less than 35 hours, and this does not take into account the fact that tapes will likely have to be changed out at the end of each step or that restoration from tape usually takes longer than backing up to tapes. If you have a contractual agreement with your customers that your data and applications will be available for them within 72 hours, your recovery strategy, barring any difficulties along the line, takes at least half of that time. In a case like this, differential backups may be a solution to your problem, at least for some of your systems.

The combination of the consideration of backup strategies to fit within the confines of the organization's requirements and the ability of those backups to meet the requirements of the organization is considered to be matching the strategy to the operational constraints. In all aspects of backup and recovery, disaster recovery, or business continuity planning, it is critical to remember the operational constraints of the organization. If any part of the plan doesn't fit within those constraints it may be functional, but it will not likely meet the needs of the organization and will therefore not be a valid solution or consideration.

Although the backups we have talked about so far can provide a solution to the recoverability of the electronic files stored by a company, the DR team must be aware that not all the company's records are stored on computer systems. There will be a certain number of records that are stored on other media, such as paper or microfiche. Planning for the storage, backup, and recoverability of these records is as important as for the computerized records.

Backing up paper documents is as simple and as labor and space intensive as making copies and shipping those copies offsite to be stored in a climate- and access-controlled location. Alternatively, an organization can choose to digitize these paper documents, having those documents scanned and stored electronically on tape, disk, CD, or DVD. Similarly, with microfiche and other non-electronic documents, duplicate copies, in one form or another, should be created and stored in a secure location.

These non-electronic documents are often overlooked in the disaster recovery effort, and are often critical pieces of information that get the organization into trouble. Older data, data from before much of the information was stored digitally, is required by governmental regulations to be available for inspection for up to 20 years or more. If an organization finds itself in need of this information and finds that it is later unavailable, it can be catastrophic.

Storage area networks (SANs) and network-attached storage (NAS) have had an impact on the backup and recovery decisions that an organization makes. Because the amount of data that can be stored on a SAN or a NAS can be significantly greater than that stored simply on internal or regular RAID array, the time required to back up these devices can be significantly greater than on common storage devices. Further, because on any given network storage device

an organization can store data from multiple heterogeneous operating systems, the backup and recovery solution needs to take these differences into account.

For more information, see the following article: **http://www.ameinfo.com/39672.html**.

Storage Backup and Recovery Tools

There are many backup and recovery tools on the market today. It is important to not tie yourself to a tool that is proprietary enough that you can't find a recovery site easily that will be able to handle your recovery needs or that you couldn't change recovery sites after a period of time if an earlier site no longer meets your needs.

Managing Stored Data and Applications Different vendors bring a variety of products to market. A simple search for data backup solutions yields links to many products that could provide you with the solutions that you are looking for. There are many others, and a fairly exhaustive search should be carried out to determine the optimal fit not only for an organization's situation but also for the environment it has as its infrastructure. Many solutions are targeted to Windows and Unix, whereas other vendors target primarily IBM-based products.

The following is a list of storage solutions that are applicable for many businesses.

- Veritas (**www.veritas.com**) provides storage protection, automation, and performance solutions for a wide variety of operating systems and databases.
- EMC (**www.emc.com**) has a full line of data management and backup and recovery solutions.
- EMC Legato (**www.legato.com**) provides data protection and availability products for an assortment of operating systems and platforms.
- Network Appliance (**www.netapp.com**) is a high-performance file server, filer, and caching system for Windows and Unix systems that provides the ability to back up and recover data in an extremely rapid manner. It also provides access to analyst reports, customer stories, data sheets, technical reports, videos online, and a substantial glossary of terms.
- Storage Tek (**www.storagetek.com**) brings with it data protection, archiving, storage productivity, and industry solutions. It also has a substantial library of case studies and technical and white papers.

There are, of course, other alternatives, and there is always the home-grown method that relies on scripts and utilities that are resident in either the operating system or that can be built and stored in the organization's software and script library.

FYI *Backing Up Scripts for Recovery*

If you are maintaining your own backup and recovery scripts, you should take care to have an external source for those scripts—one that is not reliant on the recovery of the data to the server as a means to access the scripts. It would be very difficult to recover the scripts if the scripts are needed to recover the data. A minimal backup would include a script library that is maintained independently and retrievable with just the operating system utilities such as copy (regardless of the OS there is a copy command).

The Impact of Storage Area Networks (SANs) on Recovery It is often difficult to find an elegant and cost-effective method of disaster recovery solution for an environment that uses a SAN infrastructure. SANs are being used more and more as a solution to an organization's data storage challenges.

A SAN is a network of storage disks. A SAN typically connects multiple servers, usually with heterogeneous operating systems, to a centralized pool of storage. The idea is that rather than having to manage hundreds or thousands of servers, each with their own set of storage disks, by using SAN technology you can improve overall system administration. Organizations choose this type of architecture so that they can centrally manage the storage resources with a single backup and recovery solution, a single disk maintenance solution, and a single schedule over the entire organization. This architecture can create issues, however, when it comes time to perform disaster recovery because of the heterogeneity of the operating systems and the pooled storage model.

As a pooled storage solution, the SAN doesn't care what is put on it, what the format is, what the operating system is, or whether the data is ASCII or EBCDIC. But you will have to take care that your backup and recovery solutions take into account that there may be data included from multiple operating systems and for multiple purposes.

Further, because the storage solution is likely to be used by nearly every system in the organization, additional care needs to be taken when considering the recovery timeline, the order in which the data needs to be recovered, and the criticality of the data. It may be that none of your data will be available until all of your data has been recovered, and this may not be acceptable to the business as an entity.

Determining what a cost-effective DR solution is when you have standardized to any extent on a SAN storage solution is based on your organization's needs. Some solutions are based on host-based data replication over a network (LAN or WAN) to a remote and often dissimilar SAN technology. Alternatively, storage virtualization is seeing growing acceptance and can allow for the replication of the data across the SAN or to dissimilar storage arrays.

FYI *Storage Virtualization*

Storage virtualization sounds like a complex concept, but it is simply the transparent amalgamation of multiple storage devices into what appears to the end user and the other hardware to be a single storage unit. A high-speed network is typically key to allowing for the elegant use by multiple heterogeneous operating systems and hardware solutions to access this virtual pool rapidly. Storage virtualization is typically handled by software masking the underlying hardware details.

Current Trends in Recovery In recent years, more and more organizations have realized that any event that is disruptive of their business can be a crisis. Acknowledgment of the impact of the Sarbanes-Oxley Act, the Gramm-Leach-Bliley Act, and HIPAA (the Health Insurance Portability and Accountability Act) as realities of everyday life and of nearly every business's daily life has driven many of these realizations. These kinds of regulatory issues have been the impetus for many organizations to launch their own disaster recovery planning.

FYI *Governmental Regulations*

The Gramm-Leach-Bliley Act, also known as the Financial Modernization Act of 1999, includes provisions to protect consumers' personal financial information that is held by financial institutions. (**www.ftc.gov/privacy/privacyinitiatives/glbact.html**)

The Sarbanes-Oxley Act was designed to radically change the way that publicly traded companies recorded and reported inancial information. (**www.ey.com/global/download.nsf/Russia_E/EY_Sarbanes_9_12_02e/$file/EY_Sarbanes_9_12_02e.pdf**)

HIPAA addresses health and insurance information privacy protection for individuals. (**www.hipaa.org/**)

More governmental regulations will likely come in the near future, and with these regulations will come evolving requirements for disaster recovery and implications for additional organizations.

Terrorism is becoming more widespread and is impacting more and more lives daily, but until recent events (9/11/2001 in the United States, London's transportation system bombings in 2005, and the continuing global strife over the recent years) it was not always looked on as an event that would touch many organizations. Organizations are starting to take this as a wakeup call. Even if they don't believe themselves to be at risk of being affected by a terrorist attack, they are starting to realize that they are not beyond the effects of disasters. They are also realizing that they are relying on organizations that might be affected by terrorist attacks (utility companies or transportation companies, for example) and are taking steps to mitigate these risks.

Data and information are becoming more of a critical concept for more organizations, and the demand for realistic 24/7 access to data is becoming the norm rather than the outlier in the statistical analysis of organizations. This demand for continuous access to data and information is impacting more disaster recovery plans, and service level agreements are equally impacted by the tightening timelines. This means that backing up data has to evolve to meet the need, as do recovery techniques and attempts to meet the tighter requirements.

Because email is becoming more critical as a means of business communication, and because it is becoming more relied upon as evidence in litigation, lost emails, or emails that are not retained in a sufficiently secure manner, can be a disaster recovery event. With more organizations based entirely on the Internet and with email being the primary means of communication with even a greater number organizations, the necessity of maintaining accessibility to the email system and to the historic information gathered from this email information is critical. This is true not only from a disaster recoverability perspective, but from the perspective of security as a whole. We not only have to be able to accurately recover and provide access to email information, but also be able to accurately report on and restrict users who have had access to that email and provide its relevant change history.

Restoring Communications and Recovering Users

Recovery of the business is nearly pointless unless you can restore communication with and recover all processes to the ultimate end users. Not only will this allow your business to continue its line of business and to hopefully grow and thrive, it will foster the goodwill that is necessary in doing business with both your customer and every stakeholder in the organization.

Determining Vital Users with BIA Recall that you previously have done a business impact analysis (BIA). When it comes to determining which users to restore connectivity to and in what order, this will become, yet again, an invaluable tool. One of the areas of analysis for the BIA would have included those users, both known and implied, that would be impacted by an event. Assigning weights to these users allows you to determine in what order to recover not only the software and hardware resources, but in which order to restore connectivity and accessibility to given users or given user types. It is important to restore

user access in a logical and correct order not only to critical systems and functionality but also to those areas that will be most directly impacted by service level agreements.

Rerouting Voice, Mail, and Goods Delivery So you have relocated your systems and you have relocated many of your users and business units if necessary, and you need to get back to business as usual. But to really be back to normal, communication (other than strictly electronic digital communication) also needs to be re-established. Mail (payments and invoices, naturally, but other communication as well), voice communication, and delivery of goods to alternate locations need to be taken into account.

Voice communication is a vital link in nearly any organization's business. What's more, there is no acceptable alternative to voice communication in many situations. It is the manual alternative in most cases for any other alternatives, and it remains the communication solution of choice for most people. The ability to know from the tone of voice that both parties understand what you are discussing is important. The sound of another human voice on the other end of the line to provide information or to hear out a situation is often necessary to allow all parties to maintain a comfortable relationship.

There are different techniques that allow you to minimize the chances that voice communication will be affected, including redundancy and diversity of implementation. But these don't address what to do if there is a catastrophic disruption in service. In this case, you will need to work with the telephone company to determine the best way to handle the situation. Discuss location with them *before* there is a disaster situation. If you wait until there is a disaster, you may not be the only one making similar requests, and this could cause delays. An ongoing relationship with the telephone company or companies involved will prove useful when the disaster is declared.

Rerouting voice communication can be a challenge, depending on your location and your method of communication. The simplest solution is to standardize your organization on the technology of Voice Over IP (VOIP). VOIP is a method of taking audio and analog signals, converting them to digital signals, and transmitting them, at least part of the way, over the Internet. In this way voice communication can be routed to anywhere, the IP address implies routing. Redefinition of the IP address in a routing table can mean that the voice communication can be routed to an alternative location.

Mail is much easier to reroute. Simply submitting the proper forms will allow mail to find its proper locations, and this solution is customizable enough to route mail directly to where the receiving location needs to be, even when multiple locations are involved. Again, however, a working knowledge of where to get the proper forms, or maintaining a stock of the proper forms with the disaster recovery materials, can minimize downtime.

Rerouting deliveries of goods and services may not be as simple as rerouting mail delivery. Special arrangements will need to be made for delivery, either in postponing the delivery or in rerouting deliveries during the time that you are in a temporary location. This should include not only the time you are in the new location but also the period shortly after, to take into account orders that are in process and not yet delivered at any point in time.

There may be orders that are in process at the time of disaster and already out for delivery. These orders may need to be rerouted and there may be additional cost associated with either the return of the orders or with rerouting the deliveries to an alternate location. This will be a sunk cost that won't likely be recoverable, but most suppliers will work with you.

Eliminating Network Single Points of Failure This is by far the most important thing to take into consideration when looking at backup and recovery. For every single point of failure that you can eliminate, that is one place where your system may be robust enough to make disaster recovery, regardless of the level of loss, irrelevant.

A single point of failure is any point in your system where a failure will result in loss of the system, loss of the network, or loss of ability to access information. This single point can be hardware or software, and is often not even thought of as a point where something might fail. Some single points are difficult to eliminate. Even if they can't be eliminated completely, recognizing them as points where failure can occur and taking steps to monitor and maintain these points is a way to keep them healthy and to recognize that if any issues appear to be presenting themselves they can be taken care of as quickly as possible.

Elimination of these single points of failure doesn't come without cost. You will be duplicating hardware and, often, software systems. The cost-benefit trade-off in each case needs to be considered at each point that you uncover.

But wait. Not every single point of failure is directly connected to your computer. Often the single point of failure is a user who is overly tired, under trained or who thinks that he or she knows more than they actually do. Several years ago, the I LOVE YOU virus wreaked havoc on many organizations. An accounting user at one company knew just enough to understand the language in which the virus was written but didn't realize exactly what the ramifications might be if the virus was opened. He associated the virus with Adobe Acrobat, thankfully, and called the local help desk to help him to unassociate the virus with the software so he could find another way to open it and read the code.

Just as often, a user whose IT department is less diligent in keeping computers updated with the current security patches and virus definitions can be the

cause of failure in a system or in an organization. Remember, a disaster is not necessarily the catastrophic loss that accompanies a hurricane or tornado, it can be simply the extended inability of an organization to continue with the work that needs to get done.

One way to eliminate this kind of single point of failure situation is open communication and training. Another is diligence in the maintenance of the systems connected to the system. Your system is only as robust as its least secure single point of failure, regardless of whether that point is hardware, software, or humanware.

This is just as important when recovering the organization as it is when examining the organization's home network. When you are looking at where to recover, examine where there might be single points of failure. When constructing the contracts, make sure that you read carefully and negotiate appropriately so that you eliminate as many points of failure as possible. Once you are at the disaster recovery site, there are few alternatives if there are further disasters.

Connecting End Users One of the things that is often overlooked when considering disaster recovery is where your end users will be located, and what kind of connectivity they will be able to rely upon. Although you will likely have control over the definition of IP addresses assigned to the new hardware, you may have to take into account firewall ports and VPN or dial-in access to the network.

Further, you may need to find a location to house your users. Many organizations rely on temporary trailers or set up temporary office space in conference rooms in locations not affected by the event. It is important to think through where these users will need to be, and plan for the rapid implementation of establishing these people into their temporary office space.

Not only is it important to make sure that these end users can connect as rapidly as possible so that the company can return to business as usual as rapidly as possible, it will allow your users to regain a sense of control over their lives and over their surroundings. At a time when they are likely feeling that things are out of their control, the security that you provide to them will allow them to get back to normal as rapidly as possible.

Summary

You should now understand the critical steps in assessing a system to help you best address any vulnerability in the system. It is also imperative that the person conducting the security audit document the specific steps taken as well as any flaws found, and what corrective actions were taken.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. One of the first considerations that comes into play in a disaster recovery plan is _____.
   A. data recovery
   B. data backup
   C. data recording
   D. business process

2. The quickest recovery methodology is _____, because your access times do not include the time that it takes to transfer data from another medium.
   A. tape backup
   B. floppy disk
   C. disk-to-disk
   D. portable drive

3. When determining a backup schedule, you need to take which of the following factors into account?
   A. Recovery point objective
   B. Recovery time objective
   C. Maximum allowable downtime
   D. All of the above

4. Which of the following types of failure often goes undetected for extended periods of time?
   A. Software
   B. Hardware
   C. Technical
   D. File recovery

5. Locations, other than your typical business site, where a duplicate of your data is stored and ready for access in a moment's notice are known as _____ sites.
   A. backup
   B. hot
   C. warm
   D. cold

4. Develop a general plan to recover the application as quickly as possible, which will require bringing back both the front- and back-end servers.

Project 6.2: Emergency Operations

1. A major hurricane has struck the area where your company has its central data center. There are offsite backups and hardware at three other centers around the country.
2. Assume that the area is suffering greatly from the hurricane's impact, and your company cannot continue operations until the central data center's operations resume.
3. Briefly describe the activities that need to occur to set up the emergency operation center for your company, and where your company should locate the center.

Project 6.3: Complex Vendor Emergency

1. Your company has relationships with several upstream vendors. Many of these vendors are local to your operation from a supply perspective.
2. Your company can do without components from any single vendor for up to one week but no longer, or production of your primary product will cease.
3. The executive management has asked that the company consider developing a strategy for surviving a longer outage with up to three suppliers. Develop a plan or strategy that would allow the company to survive for up to two weeks without supplies from three of the primary vendors.

Case Study

You are helping to prepare the recovery plan for a medium-sized business in the manufacturing industry that has one physical location. All operations occur at this location. Construct a recovery charter for this company.

Chapter 7

Developing Procedures for Special Circumstances

Chapter Objectives

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

- Identify emergency situations that may occur during a recovery.
- Determine what can be done if an emergency occurs during an emergency situation.
- Assess the risks associated with disaster recovery.
- Identify gaps in emergency recovery situations and plan accordingly.

Introduction

The best-laid plans often go awry. What will happen when all the hard work you have done suddenly hits a snag and you have to deal with emergencies during the emergencies? We have seen that it is important to have a plan for just about any contingency that might arise in an organization. If the primary plan cannot be followed, having a backup plan will assist with alleviating the stress that will occur. This will also provide you and your organization with a fall back point to assist those performing the recovery and provide stability for them when they will likely need it most.

The disaster recovery effort, whether a planned practice or an actual emergency situation, is stressful enough. Change is difficult for many people, and the added change of changing the changes may be too much for them in such a short time. Knowing that there is a plan for this kind of situation will ease the peace of mind of not only those involved, but also those stakeholders who are relying on the recovery to occur seamlessly.

Emergencies During the Emergency

Planning for emergencies is something that a well-prepared organization does. But what happens when—even if you have a disaster recovery plan, have tested it, and are comfortable with the plan and its execution—the unimaginable occurs: you have to declare a disaster and go through the recovery plan for real? Do you think that is the worst that could happen? It isn't.

Imagine that you have declared your disaster and are proceeding with the recovery. Something happens during the recovery that constitutes an emergency or yet another disaster. What can happen now?

Does your organization's preparedness plan encompass the potential eventuality of having to face an emergency at the disaster recovery site? Have you prepared for such an eventuality? Have the necessary steps been taken to ensure that you could pick up and move to yet another site in order to recover your organization? What will this do to the SLAs that your organization has in place? Does your recovery site have plans in place for these kinds of eventualities?

FYI *Wording SLAs*

It is possible to word SLAs in such a way as to make sure that the organization is covered in the eventuality that an emergency within an emergency should present itself. It is in an organization's best interest to make sure that the SLAs that are in place are written such that, should an emergency occur during the recovery effort, the clock is reset to that time and the SLAs are not violated.

Concrete, realistic, and measurable commitments are important for a successful SLA. If your SLA includes a clause identifying the potential for an intra-emergency emergency, specify that your organization will have the additional days necessary to relocate and re-recover the data. If you have agreed upon 72 hours for recovery, word your additional emergency clause to specify that, in the event that a disaster that is declared within a disaster, an additional 72 hours will be allotted for the recovery from the second disaster and that this 72 hours will commence at the declaration of the disaster and will supersede the time allotted for the declaration of the original disaster.

Not limiting the extra disasters to a second one will also cover any eventuality that may occur that causes the organization to declare disaster after disaster. Although this is definitely a worst case scenario, it should not be discounted.

Support Contracts

Chapter 2 looked at the necessity for support contracts for your organization as well as for a disaster situation. But what about support contracts as they might apply should a second emergency or disaster occur while the organization is in the process of recovery from the initial disaster? Does the organization have the contracts in line to be able to cover itself should this kind of situation arise? Arrangements with many software vendors can put the organization on better footing should the need arise to make this kind of adjustment in recovery.

Further, it is important that the organization have support contracts with the recovery location should it become necessary to relocate the disaster recovery effort to another location. Although it may seem overly cautious to talk to the recovery site company about the eventuality of the recovery site either not being available when needed by the organization or needing to relocate to an alternative site in the midst of an ongoing recovery, it is in the best interests of both the organization and the recovery site to have all of the relevant information on the table up front.

Because of the nature of its business, the recovery company should already have in place arrangements for this kind of eventuality. If it does not it should raise questions about the suitability of the recovery site for the choice of recovery location. Because the organization doing the planning should have any emergency eventuality uppermost in its mind when determining what contracts to sign and what recommendations to make, the question of what would happen should an emergency befall the recovery site should be among the questions asked. However, often these questions are not asked, and the contracts that should be in place are not.

FYI *Home-Based Businesses*

Although this section appears to be directed primarily at organizations of significant size, it is just as important for small organizations, including home-based businesses, to be aware of the need for plans for what might happen should additional emergency situations occur.

Many of the needs of this type of organization are commodity in nature, and they can be acquired in almost any location with very little notice. However, many cottage industries rely on either customized software or software that is not always freely accessible in all locations. Special arrangements may need to be made with software, hardware, or machine vendors to speed shipment of specialized tools to alternative locations. If those alternative locations should have to change, it is important that the shipments be re-routable with little burden on either the manufacturer (reseller) or the buyer.

Disaster Recovery Contracts

As we have already discussed to a small degree, it is important for all concerned to have disaster recovery contracts and contracts concerning what to do in the eventuality of an emergency occurring in the midst of an emergency. But what happens once you have your contracts? Organizations typically keep these documents in safe locations, locked up where they will not be affected by emergency situations or where they can be easily accessed in the event that they need to be referenced or renegotiated. These precautions are all well and good, but how accessible will these documents be in the midst of an emergency? If they are needed for reference at any point, are they in a location or format that is easily accessible to the recovery team? If it becomes necessary to move the recovery location, can the necessary documents be easily and readily accessed?

It is important that all relevant documents, or more appropriately either electronic or photocopies of these documents, be kept with every set of backups that the organization maintains for the eventuality of a disaster recovery.

IN PRACTICE: BS Sunlight Application Service Provider

BS Sunlight is an application service provider to many organizations. It has in place a disaster recovery plan and tests it every year to some degree. It has a backup strategy that is commensurate with its recovery plan, which covers all SLAs.

The backup strategy dictates that a full set of backup tapes be maintained in waterproof cases and stored offsite in a secure location, and that these backups can be obtained only by a certain set of management members and only under certain circumstances.

The waterproof cases contain a full set of all manuals on CDs and DVDs for all of the systems that will be necessary to recover. The recovery documents—in hard copy, on CD and DVD, and embedded within the backups of the servers—are included, as are copies of all contracts that may become necessary, complete with full contact information for all vendors and all support information in both hard copy and scanned or stored on CD and DVD format.

Because BS Sunlight has determined that it wants to be able to recover the business regardless of the severity of the emergency situation and regardless of the situation in which it finds itself, it has made the business decision to take these steps so that it can recover the business regardless of what may befall either the primary site or one of the backup sites with which it has contracted.

Preparations

When an organization prepares for the eventuality of an emergency in its primary location, it takes certain steps in order to become prepared. Backups are put in place and are taken routinely. Security measures are put in place and followed. But when individuals are thrust into unfamiliar surroundings or are working under conditions of additional stress, shortcuts are often taken, either deliberately or inadvertently, that could cause the organization to be at risk.

Under these circumstances, backup are often not undertaken at as regular an interval as they were in the primary location. There are even organizations who have made the decision to outsource their IT services to better fit with their business plan yet who assign their backup tape rotation to one or more employees whose main job (due primarily or in part to the outsourcing) is not IT related. The employees are supposed to care for the tapes and make sure that they stay safe and that the logical rotation is maintained. Because the employees do not understand the ramifications of what is being asked, backup jobs may be missed, tape rotation may not occur in a timely manner, tapes may not be maintained as diligently as necessary, or rotation may not occur in the correct order and the necessary backup may be overwritten before it should have been. It is often necessary to make sure that the persons assigned to the backups are either IT trained or have sufficient understanding of the steps necessary and the criticality of the job before they are assigned to perform the tasks. Even well-written documentation, if there is not underlying understanding at some level, may not be enough.

Identifying the Gaps in Your Recovery Plans

In any organization's attempts—particularly early attempts—at disaster recovery planning and testing, gaps in what people think will happen and what actually happens appear. One prime example occurred in the wake of the September 11th terrorist attacks. Many organizations realized that their disaster recovery plans were lacking and decided that they needed to do something about it. The gaps that many organizations found were made apparent because of their inability to recover in a timely manner from the disaster.

When gaps were discovered the companies worked hard to overcome them, and nearly 2 years later, when the August 2003 blackout hit the East Coast of the United States, many organizations had installed electrical generators that would provide their computer systems with power. This allowed them to continue to do business on a limited scale. Even companies not directly affected by the terrorist attacks worked on making their DR plans

more robust and were better prepared for the blackouts that crippled much of the Northeast for more than a day.

Many of the gaps that organizations strove to fill in this example were not gaps that were found in the process of testing efforts or even in the midst of their own disaster recovery. They were, however, still gaps in planning.

In the cycle of planning, backing up, testing recovery, and restarting the cycle, organizations typically find their own gaps, such as places where backups have been insufficient, where systems were missed entirely in the inventory phases, and where the knowledge and ability of those involved in the exercise is found to need augmenting. These are discussed next.

Backups

Backups are one of the first areas where gaps are identified. Either the backup is not broad enough to capture enough of the system to be able to recover in a disaster situation, or the backup media and methods are not entirely compatible with the hardware and topology of the recovery location.

IN PRACTICE: BS Sunlight's Backup Strategy

BS Sunlight put its disaster recovery plans in place. Backups were taken on a regular basis and testing was planned for the summer months.

The disaster recovery team made the journey to the recovery site and started the recovery drill as if it were a disaster. The recovery media made it to the recovery site successfully and on time. Most of the systems were up and running with sufficient time to make sure that the SLAs would not have been missed had this been an actual disaster recovery situation.

However, on one set of systems, it was discovered that the hardware that the recovery center had substituted in place of the hardware that the organization was using (the recovery center substituted newer models for older, which was within the bounds of the contract) was not completely compatible with the backup media or the backup method chosen by the organization. This meant that several key systems in the organization were unable to be recovered during the testing. Both the methods and media used for the backups and the contracts with the recovery location company were examined in an effort to make sure that this gap did not continue and was accounted for in subsequent tests.

Testing

Gaps are often discovered in testing efforts. These typically are not found during the testing session in which the gaps occur. They often are discovered in subsequent testing when tests fail, or in actual disaster situations when systems that are turned over to the end user do not perform adequately.

These gaps normally can be uncovered if the test is run as if it were an actual disaster situation. In this case, the actual systems at the primary location are left intact and running, but the recovery effort progresses as if it were a disaster. Often, those involved in the recovery effort are not entirely sure if it is a test or a recovery.

When the recovery is complete, the entire system is switched from the primary location to the backup location and business continues for a period of time on the recovered system as if it were the primary location. In this way, more of the system can be tested and more gaps in testing are likely to be discovered.

Systems

It may seem counterintuitive that an organization can overlook critical systems in its planning and recovery efforts. However, it is not uncommon for an organization to find that, either in the midst of a test or, worse, in the midst of a recovery, one or more functional areas of the organization have overlooked a PC or server sitting under someone's desk that is business critical to some process.

Typically these oversights are honest mistakes. A department will have acquired a system that runs only on a platform that the organization does not have the expertise to maintain. They don't think about the system unless there is a problem with it, and then only to the extent necessary to keep the system functioning.

Backups are often not done on a regular basis. If they are done, they may not be tested fully. Further, accurate records on license and support details may not be readily available.

Because functional people often consider that they only work on "the accounting system" or the "HR system" they don't (and maybe shouldn't have to) delineate the difference between using the primarily supported systems and the ones that are supported outside official administration circles.

Hopefully these missing pieces are found in the testing phase of the disaster recovery planning. If they are not uncovered until a true disaster recovery there may be little recourse for the organization other than to attempt to, as rapidly as possible, recreate the system to whatever extent possible and attempt to recover the functionality, possibly with limited or no access to data that might be associated with the system. At best, the organization may see the recovery take longer than expected. At worst, it may see an extended impact on the ability to do business for a longer period of time than should be possible otherwise.

If the missing system is identified in the testing situation, extensive notes should be made on what the gap is and what systems are affected by the missing system and its functionality. When the trial recovery is over, special effort should

be made to make sure that the system in question complies with the organizational standards on backups and recoverability. At the very least, the contact information and the knowledge of the system should make it into the recovery documentation.

People

Some of the biggest gaps found in disaster recovery planning concern the people involved. People gaps are discovered in not only the area of people involved in the disaster recovery (or the recovery test) but in the areas of planning and end-user testing.

If the wrong people are involved in the planning effort, other gaps can be discovered because of the missing knowledge of systems that need to be included or in the ways that backups need to be accomplished to meet the need of recovery. If all levels of management are not included, or all stakeholders in the planning process are not taken into account, the potential of not meeting SLAs might prove to be costly to the organization.

People gaps that occur in the disaster recovery testing effort, particularly in early testing efforts, can be equally as costly. Although eventually the organization will want to have as many people as possible involved so that they can be considered trained in the recovery process, in early testing attempts it is important that knowledgeable people be involved because these tests will provide the foundation for all future testing and for the ultimate test, an actual recovery.

In later testing efforts, it is often beneficial to have inexperienced individuals placed into roles that they may not be well suited for. This will show where there are inconsistencies in the plan or in the documentation. Once the plan is established and the tests have been successful for a period of time, it is important to foolproof the tests so that anyone that needs to be drafted to fill positions can follow the documentation successfully. This can give the organization the assurance that anyone can be called upon in a disaster to fill the position and provide the needed support.

Disaster recovery professionals often suggest that you write your disaster recovery instructions so that anyone (even employees with no ties to the IT department or even the day-to-day workings of the company) could use them to recover the system. If this is the case, you can rest assured that your documentation is sufficiently thorough to provide the necessary coverage. It is not a comforting thought that most organizations are not sufficiently well prepared enough to make this claim.

You may have all the right people involved in the planning and recovery efforts, but if you don't have the right people involved in testing that the recovery is successful and accurate, then all of the effort that has been put into the planning, the backups, and the recovery will be for naught.

Again, finding out that you don't have the right people in place isn't the worst thing that can happen, particularly if you find out before a disaster has been declared. First, you have gained the insight of the people that you do

have, right or wrong, and no insight will have gone to waste. Further, you can augment the team, wherever the inconsistencies are, with more appropriate people. It would not be the recommended way to staff the planning or the recovery testing team, but experience is gained from the existence of the inconsistencies.

Identifying Disaster Recovery Risks

There are always risks associated with doing business, even in a perfect world. There are additional risks associated with being in a disaster recovery situation.

Location

One of the risks associated with a disaster recovery situation involves the fact that everyone is in an unfamiliar location. Although this is not typically considered a risk in itself, it contributes to risk because of the heightened stress associated even with a testing situation. It can have one of two effects. Often those involved will become overly cautious in their actions and not perform as rapidly or as efficiently as necessary. This can be even more of an issue if everyone is aware that the drill that they are involved in is not an actual disaster situation but a testing situation. Conversely, the individuals can become less cautious than they may be in other situations, more aggressive, and pay less attention to detail than might be prudent. In this case, mistakes can be made. These can be costly if the drill is not a testing situation; they could mean the difference between the company meeting its SLAs and having to pay fines for not meeting them.

If the individuals involved in the recovery are aware that it is a test and not a true disaster, the exercise may be viewed as a paid working holiday, and receive less seriousness than should be accorded. When this is the situation often less is accomplished than should be. This is one reason that some organizations, when planning tests, don't include everyone in the planning, but spring the trip on them under the guise of an actual disaster. This means that fewer people will have the tendency to look on the situation as relaxed free time and will more likely strive to meet the deliverable timelines.

Situation

Because the recovery team is in an unfamiliar location and working in an unfamiliar environment, they are likely to be less mindful of security precautions if they are different than those at the primary work location. Security measures often are more stringent at a recovery site, and people accustomed to more lax security measures may inadvertently counter the security measures that are in place.

Further, old work habits and working hours may remain. Disaster situations, regardless of whether tests or declared disasters, call for extraordinary practices—longer than usual work hours, altered work schedules, and working under conditions other than what everyone is used to. Unaccustomed managers will deliver instructions and directions and unusual work environments will make many of those involved uncomfortable.

Systems

There are additional risks in that recovery will be performed on different systems than everyone is used to dealing with. This may not seem like an apparent risk—after all, a system is a system. It is little more than a bunch of commodity pieces put together to form more commodity pieces on which the recovery team will recover the organization's software and systems. However, every system has its own idiosyncrasies and its own security precautions in place to keep intruders out.

It is common to assume that because there are contracts in place, the recovery will go smoothly because the hardware is compatible. Hardware is often considered to be interchangeable because it is a commodity item. This is not necessarily the case, however. There are occasions when, even though contracts are in place and hardware has been provided as compatible, it is difficult, if not impossible, to recover.

IN PRACTICE: BS Sunlight

The systems administrator and database administrator for BS Sunlight run a regularly scheduled recovery test, which proceeds in the usual manner. However, for this recovery testing, the hardware (which the contract states must be compatible with what the organization has) was changed because the hardware that BS Sunlight used for tape backups was outdated and newer hardware, capable of reading the same kinds of tapes, was substituted.

No one, even after being alerted to the substitution, anticipated that it would be an issue. However, when the recovery started, it quickly became apparent that the new tape drives were in fact not able to read the tapes brought for the recovery. Despite the best efforts of the organization's contracts with the recovery site and the software support contracts with all of the software vendors, it was impossible to recover the operating system to such a state that it was able to be used, let alone any of the databases that

CONTINUED ON NEXT PAGE

CONTINUED

resided on that system. The gap was not in the people or in having the necessary contracts in place. The gap was that the recovery site did not have a proven method for assuring that the contracted hardware was provided and in place when needed, and that any contracted organization could successfully recover its systems should the need arise.

Summary

Emergencies can happen at any time, even when you're already in the midst of one. This chapter demonstrates the importance of preparing your organization for an emergency situation, or even a disaster situation, in the midst of an existing emergency. These emergencies may be similar to the existing emergency, or they can arise because insufficient or misdirected testing or other preparedness issues. This chapter should help you identify potential issues and protect yourself from further harm.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Having a(n) _____ plan for the eventuality that the primary plan cannot be followed will assist in alleviating the stress that occurs during a disaster.
   A. fallback
   B. operational
   C. backup
   D. functional

2. Knowing that there is a secondary plan when the primary plan cannot be followed will ease the minds of _____ who are relying on the recovery to occur seamlessly.
   A. stakeholders
   B. employees
   C. company officials
   D. all of the above

Exercise 8.2: Project Strategies

1. As part of the disaster recovery planning at a medium-sized business, you have been asked to develop a project plan to test the backups of production systems.
2. Develop an outline of the project plan for the testing.

Exercise 8.3: Change Control

1. Disaster recovery planning requires change management of the plan, as do the items involved in the recovery, since technological environments change over time.
2. In a well-supported essay, describe how you should attempt to incorporate the changes that occur to a large enterprise network over time in disaster recovery planning.

Exercise 8.4: Testing Methods

1. Sometimes it is not feasible to conduct a complete real-world test of a recovery.
2. Describe a recovery that would not be feasible to test with a real-world simulation.

PROJECTS

Project 8.1: Survey

1. Using a list of companies in your area, ask five of them about the testing of data backups or some other recovery area that needs testing.
2. Develop a comparison of how these companies approach the recovery testing and make a recommendation on how each company might improve its testing process.

Case Study

Familiarize yourself with some of the technology used by an online technology vendor, such as Dell. Write a thorough discussion of how the vendor should approach testing the recovery of some aspect of their technology with respect to the online ordering process.

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

>> CONTINUED

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.

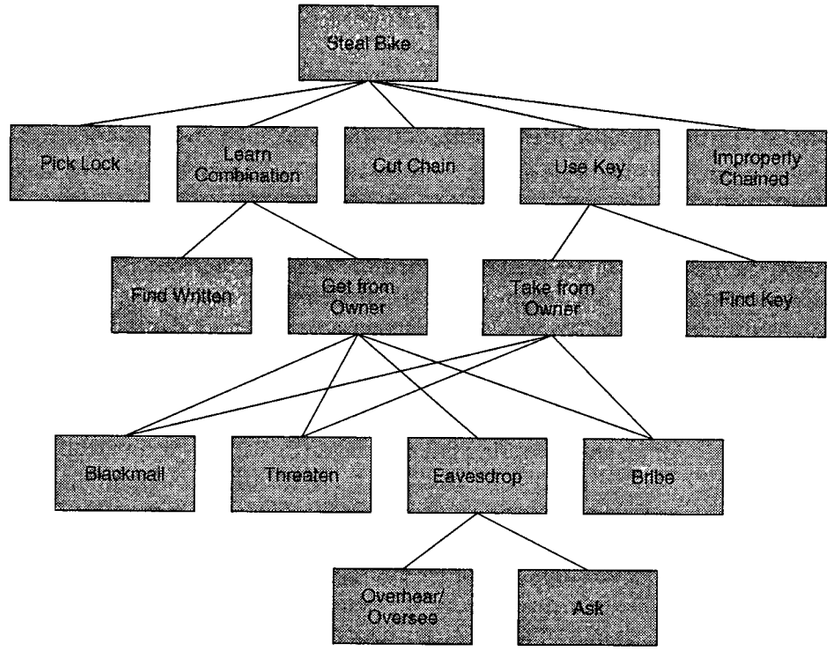

FIGURE 9.1. Threat tree for stealing a bicycle.

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

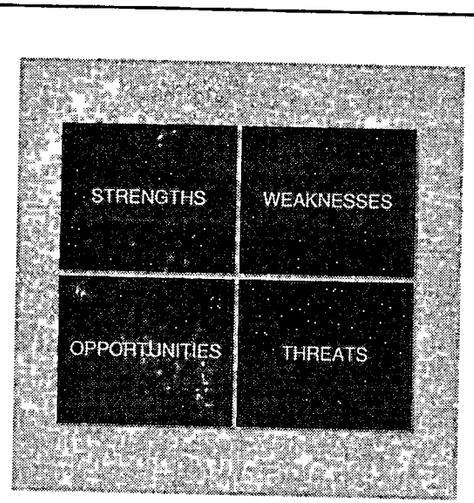

FIGURE 9.2. SWOT analysis graphic.

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.