DATA

Your Most Irreplaceable Asset

640K ought to be enough for anybody.

—BILL GATES, cofounder of Microsoft Corporation, 1981

INTRODUCTION

Most of what you lose in a disaster is relatively easy to replace. Buildings can be rebuilt or new offices leased, furniture is easily replaced, and even new computers can be purchased at the click of a button. What is not easy to replace is your competitive advantage, which is stored in the files and databases within your computer systems. This critical information is in accounting files, customer lists, part lists, and manufacturing drawings. This information is unique to your company; it is what makes your company special to your vendors and customers. It is the very essence of your company. Unlike physical assets, this information is difficult, if not impossible, to re-create once it is gone.

There are two types of risks to the infrastructure that supports your data assets: (1) physical loss due to a device failure or a disaster at your location and (2) logical loss caused by an application or user error. Physical loss is the less likely of the two, but it is potentially the most damaging. It includes incidents such as a hard disk failure, server failure, or an environmental disaster such as a fire or flood. It can affect just a single device or your entire location. Physical loss accounts for approximately 20 percent of all incidents affecting information technology resources. In contrast, logical loss includes incidents such as application errors, user errors, or a security breach. Logical failures account for approximately 80 percent of all incidents. A logical failure can be easier to repair, but it may also not be noticed for some time.

COMPONENTS OF AN INFORMATION TECHNOLOGY INFRASTRUCTURE

A modern corporate computing environment consists of components that build on each other to support the functions of the business. You must understand each of these components and how they relate to your business process to create an effective recovery strategy. At the foundation of this infrastructure are data. Figure 20-1 shows the typical components of an information technology (IT) infrastructure.

Each layer builds on the layer below, building up to the application that the user sees. The applications interact in varying degrees depending on the requirements of the organization. But no matter what the specific architecture, the foundation is the data stored on various media somewhere within the organization.

RISK ASSESSMENT

Your data is susceptible to loss or damage or both from several sources. Some key causes of data loss include:

- ***Viruses.*** These malicious programs can get into your system at any time and strike when you least expect it. Once you're infected, the virus can spread from system to system, destroying data along the way.
- ***Natural Disasters.*** Fire, flood, and high winds can all cause physical damage to systems and make your data unavailable or unreadable.
- ***Human-Created Outages.*** Systems can be damaged by a sudden loss of power, or worse yet, a small part of a data stream can be lost, causing damage that may not be readily apparent.
- ***Hard Drive Crash.*** It's not *if* a hard drive will fail, but *when*. A hard drive is most likely to fail within 90 days of being placed in service and after about three years of average use (see Figure 20-2).
- ***Laptop or Smartphone Loss or Theft.*** The value of the data stored on a laptop or other portable device usually far exceeds the cost of replacing the hardware.

Market research firm IDC estimates that approximately 60 percent of all corporate data resides on laptop and desktop PCs.

- ***Software Failures.*** Operating systems and storage area network software can fail, corrupting existing data.
- ***Application Failures.*** Applications are not guaranteed to be bug free; a bug in an application can cause incomplete or incorrectly formatted or calculated data to be written into your files.

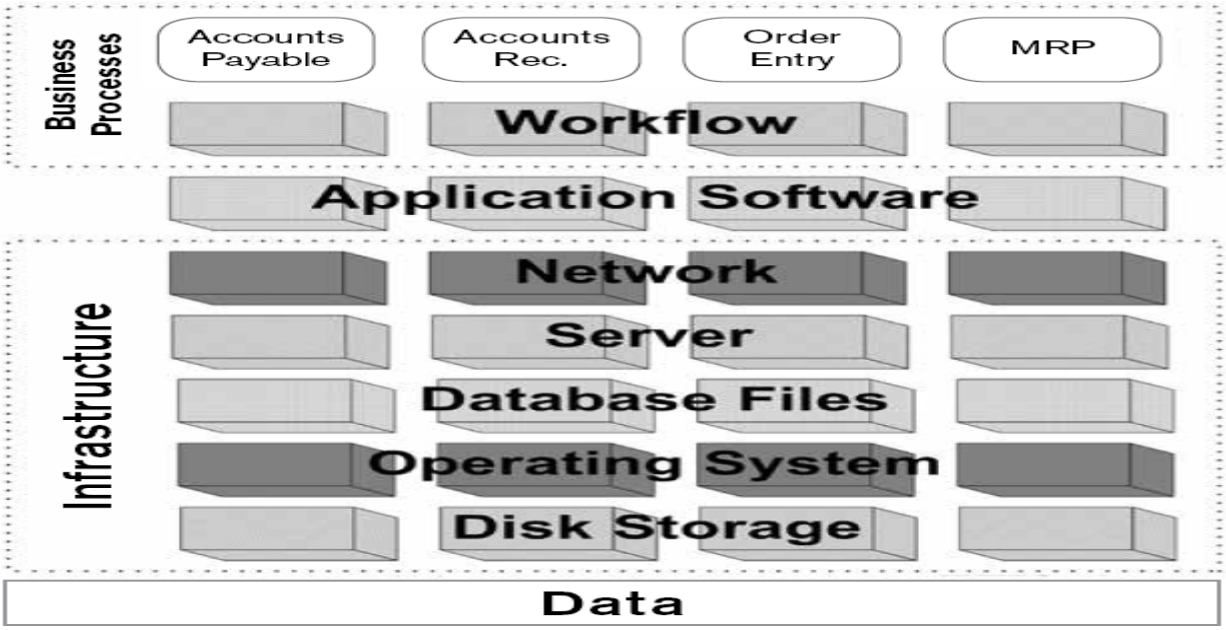

FIGURE 20-1. *Information technology (IT) infrastructure.*

Vendor Failure. If you are hosting e-commerce or other applications with a cloud vendor, your data could be at risk if the vendor suddenly goes out of business.

Approximately 20 percent of application downtime is attributed to a "disaster." The breakdown of causes:

Application code failures: 40 percent
Operator error: 40 percent
System/environment failure or other disasters: 20 percent

Source: Legato Systems

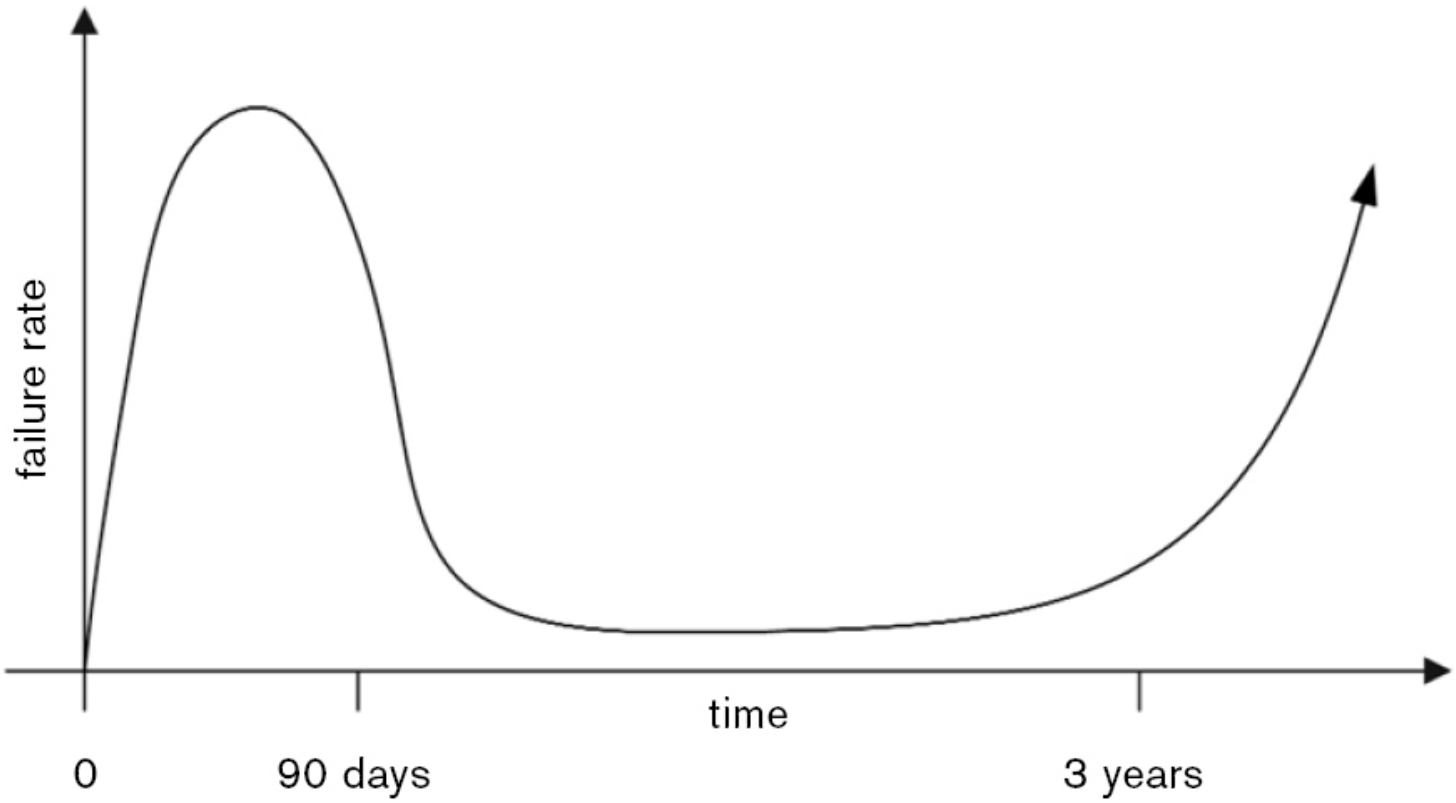

FIGURE 20-2. *Hard drive failure rates.*

There are both tactical and strategic issues surrounding the loss of critical corporate information. Tactical issues include:

Compromised Information. Your valuable information could fall into the hands of competitors if stolen by hackers or if you lose a portable device. Your competitors having this information could be more damaging than if it were simply destroyed.

- ***Lost Productivity.*** Re-creating lost data can be very expensive, especially if it must be re-created from paper records.
- ***Employee Downtime.*** Employees need their information to do their jobs; this includes people in the Sales, Customer Service, and Accounting departments.
- ***Loss of Customer Information.*** Loss of important customer records can seriously hinder your ability to serve your customers.
- ***Increased Help Desk Support Required.*** Not only might your help desk people be needed to help restore your data, but they will be bombarded by users requesting assistance and information.

In 2015, Anthem Inc. paid $115 million to settle lawsuits as a result of a data breach that affected nearly 80 million individuals, the largest cyberattack suffered by a company in the health care sector to date.

Strategic issues surrounding data loss are those that have an impact on some critical operation within your business processes. This might include:

- ***Loss of Opportunity.*** Without up-to-date and accurate information about your customers and your company, data loss can result in lost sales. If you don't have accurate inventory information, customers may order from someone else who can guarantee delivery from stock. Follow-up calls to customers might be missed if your customer relationship management (CRM) data is lost; this may also result in lost sales. Future sales could also be in jeopardy.
- ***Decreased Operational Efficiency.*** The lack of access to data will result in a greater reliance on manual processes, which will drastically decrease your operation efficiency.
- ***Inability to Support Customers.*** Without access to customer data, you will have a difficult time supporting your customers or will incur unnecessary costs providing support to which they are not entitled.
- ***Increased Systems Costs.*** Your total cost of ownership (TCO) will increase, making it more difficult to make money if margins are thin.
- ***Noncompliance Issues.*** Without accurate data, you might not be able to prove compliance with government mandates, resulting in fines and legal fees.

Other costs you may incur from a serious data loss incident include:

- ***Customer Notification.*** Many states now require that companies notify all customers potentially affected by a data breach.
- ***Litigation Expenses.*** Lawsuits resulting from a data loss incident can be very expensive.
- ***Internal Investigations.*** Time and resources will be required to clean up after a data breach.
- ***Forensic Experts.*** You may need to hire outside forensic experts to help identify any existing security weaknesses.
- ***Software Updates***. In many cases numerous software updates may be required to patch security holes.
- ***Subpoenas by Government Authorities.*** You may be required to respond to subpoenas from state attorneys general or the Federal Trade Commission.
- ***Stock Price.*** If you are a public company, your stock price may go down after a data breach becomes news.
- ***Reputation.*** Data breaches affecting credit card information can be especially damaging to a company's reputation with its customers.

CREATING YOUR DATA RECOVERY PLAN

Just like any other project, there are several distinct steps required to develop your plan to successfully recover your data after a disaster. The recommended steps, as shown in Figure 20-3, are:

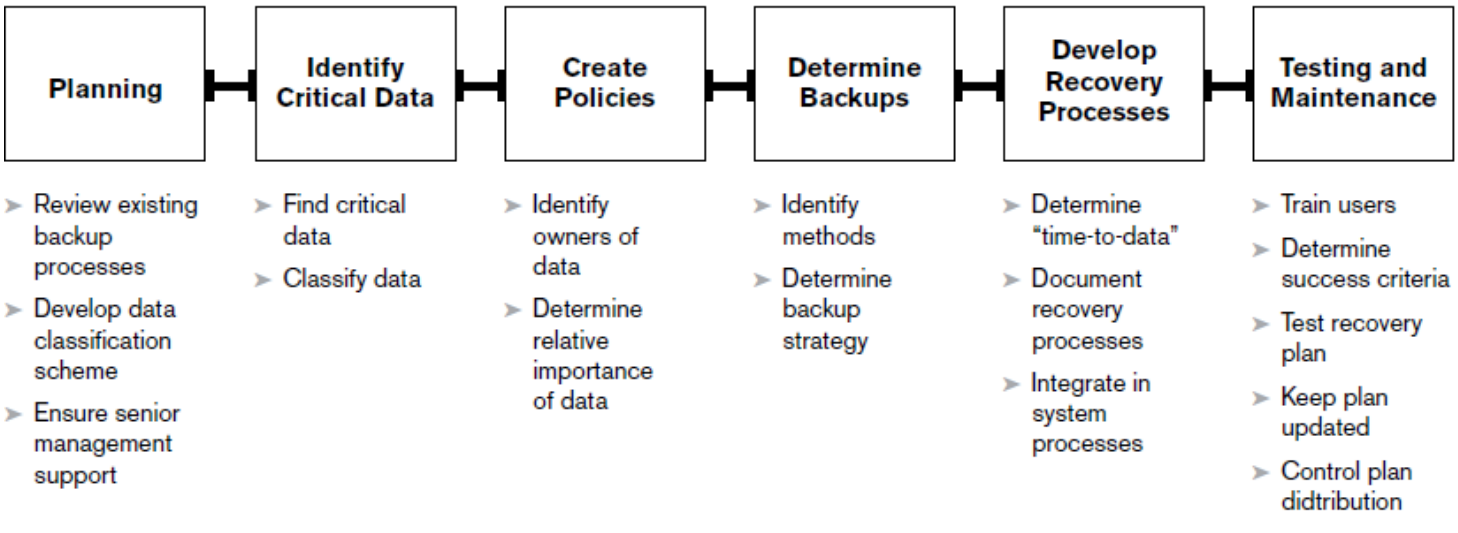

FIGURE 20-3. *Data recovery steps.*

Planning.

Identify critical data.

Create appropriate policies and procedures.

Determine type of backups.

Develop recovery processes.

Plan testing and maintenance.

PLANNING

As with any project, a successful data recovery plan begins with proper planning. Your first step should be to review data recovery expectations with key stakeholders. Find out what their business needs are, and if there are regulatory requirements about which they are concerned. Few organizations have not done any data recovery planning, so your next step should be to review the existing backup and restoration strategies. Find out what is currently being backed up and how often. Are there procedures in place to periodically test the backups? How are the backups transported and stored? What new systems have come online since the backup documentation was last updated? Are old files being backed up that could be archived and removed from the live systems?

Begin researching the most efficient and effective ways to store your backed-up data. Do you have multiple facilities that can store each other's data? So the data being stored cannot be destroyed in the same disaster, make sure the facilities are at least 20 miles apart.

You must also plan for an analysis and classification of data. What is the importance to the firm of each file being backed up? The cost of protecting the data should be proportional to the value of the data. You don't want to spend a lot of time and resources protecting data that can be easily restored by other means.

There are numerous strategies for backing up and restoring of data. They include traditional offline storage methods, such as hardcopy printouts, magnetic tape, CD-ROM, portable hard drives, and microfiche. However, online methods, which include disk mirroring, storage area networks, and Internet backups, allow for faster restoration of your data. The evaluation and selection of the appropriate strategies are critical to the success of your recovery plan. Other things to consider are whether your backup hardware and software tools are the latest versions from the manufacturer. Many firms have had a disaster only to discover that the software needed to read their backup media was outdated and difficult and expensive to obtain, or simply no longer available.

Where your data will be stored is also an important consideration. It's most convenient if your firm has multiple locations that can store each other's data. You might also have a reciprocal agreement with another noncompeting firm to store data for each other. Of course, you need to be reasonably sure that both locations won't be affected by the same disaster. You'll also need to make sure that both locations can handle the extra workload if one site is down. This option is difficult to manage and does not always work well in practice.

Another option is to use a commercial storage company that will have an environmentally controlled facility to protect the integrity of your media. It will also have tested procedures for storing and retrieving data in an emergency and can offer advice on your disaster recovery plans.

If you lose not only the data but the hardware on which it is stored, you will also need a place to set up replacement hardware. One option is to have a contract with a vendor to have an off-site facility ready if your location experiences an incident. There are several basic types of remote sites:

- ***Cold Site.*** A cold site is simply a place to store your data. It should have adequate space and infrastructure (power, communications, and environmental controls) to support your systems. This is the least expensive option, but it requires the most time to get up and running in the event of a disaster.

- ***Warm Site.*** A warm site has systems and communications ready to go, but it requires that data to be restored onto them before they are ready to use.

- ***Hot Site.*** A hot site is an active duplicate of your live systems, with both systems and data ready to go at a moment's notice. Hot sites are usually staffed 24 hours a day, 7 days a week, and are prepared for immediate action if an incident occurs.

- ***Cloud Backup.*** Your data and virtual images of your critical servers can be duplicated in the cloud. Data can be backed up in real time, and access can be switched over very quickly. Data stored on your workstations and mobile devices can also be backed up to a cloud provider, to be accessed from anywhere if a disaster occurs.

- ***Mobile Site.*** A mobile site is a self-contained transportable office custom fitted with IT and communications equipment. It is usually transported by truck and can be set up at the desired location. The mobile site needs to be configured before it is needed for it to be a viable recovery solution. If using an outside vendor, a service-level agreement is necessary to make sure the vendor is committed to meeting your needs in an emergency.

- ***Mirrored Site.*** A mirrored site is an exact duplicate of your production site, with data stored there in real time. This is the quickest way to get your business back up and running, but it is also the most expensive.

The different recovery site options offer different cost and recovery time tradeoffs. Your organization's restore time requirements and the results of your Business Impact Analysis will determine which option you choose. Figure 20-4 compares the resource requirements for the different recovery site options.

Restoring data in the fastest time possible will minimize the revenue loss caused by damaged or lost data. "Time-to-data" is a critical metric to evaluate when creating your recovery plan and is defined as how much time it takes for your users to have access to their data after a disaster occurs.

Rapid "time-to-data" is fundamental in achieving reduced downtime and maximizing productivity and system I/O rates.

Source: Strategic Research Corporation

Asset management is an important key to recovering your systems. You'll need an accurate and complete hardware and software inventory list. You'll need to know when and where it was purchased, as well as the warranty status. You'll need to know where the hardware was located and how it was configured. Your original software licenses will be necessary to facilitate getting new media from your software vendors. You'll also want to research what the vendor's policy is in a disaster situation. You'll want to know what to do if you need to quickly obtain replacement software.

IDENTIFY CRITICAL DATA

The first problem you'll face in creating your data recovery plan is *finding* the data. The amount of data being produced by business today is growing rapidly. These are not just data stored in traditional databases, but also include graphics, word processing files, spreadsheets, sound clips, and other enhanced forms of data. In many organizations, terabyte (approximately one trillion bytes) databases are becoming common; petabyte (1,024 terabytes) databases are right around the corner. And, of course, paper is still an important repository of data; these paper files are stored in file cabinets and desk drawers. Microfilm and microfiche are also still used in many organizations. For data stored electronically, there are products available for automatically discovering files and databases throughout your network

Type of Site	Cost	Equipment	Communica-tions	Setup Time	Location
Cold	Low	None	None	Long	Fixed
Warm	Medium	Partial	Partial to Full	Medium	Fixed
Hot	Medium to High	Complete	Full	Short	Fixed
Cloud	Medium to High	Partial	Complete	Varies	Mobile
Mobile	Medium to High	Servers Only	Varies	Varies	Mobile
Mirrored	High	Complete	Complete	None	Fixed

FIGURE 20-4. *Recovery site selection criteria.*

The next issue after you have found the data is *categorizing* the data. Like paper files, much of the electronic data that you create is never referenced again. You'll need to identify the critical data required to restore critical business operations. Don't forget to review ancillary data and documentation and data that must be preserved due to legal requirements.

NONESSENTIAL DATA

Much of what is stored on your file servers by users is data that is not essential to the operation of the business. This includes space-wasting data such as email attachments, Internet cache files, and personal files such as digital pictures. This nonessential data can add to the cost of backup and recovery in many ways. If you have a hot-site facility, it will require more disk storage space. If you are performing backups using tapes or CDs, additional media will be required for backups. If you are using replication to a remote location, additional bandwidth may be required to support the transfer of all these files.

A place to start in reducing the volume of unneeded files is to have policies in place that prohibit the storage of personal files on company servers. Strict enforcement of these policies can dramatically reduce the amount of data that must be backed up. You should also consider limiting the amount of storage space available to each user, which will force them to consider carefully what to store in their personal folders.

CREATE APPROPRIATE POLICIES AND PROCEDURES

Most companies do not have policies and procedures for storing and classifying data. And many that do have policies do a poor job of enforcement. Having policies that aren't enforced can create a false sense of security, which can be worse than having no policies at all.

The first step in creating policies for storing and classifying data is to identify the owners of information. All data in the company should have an identified owner who is responsible for understanding the importance and use of the data.

Once the owners of the data have been identified, develop a policy for determining the relative importance of data. You can then develop an information classification scheme. Some categories you might use include business critical, sensitive, legally required, and noncritical.

- ***Business Critical.*** Data you must have to run your business may include customer lists, production drawings, and accounting files.
- ***Sensitive.*** Data that you would not want your competitors to see might include customer lists, employee lists, and production process documentation.
- ***Legally Required.*** This is information that you need for compliance with government regulations, such as the Occupational Safety and Health Administration (OSHA) compliance data, Environmental Protection Agency (EPA) information, and hiring data.
- ***Noncritical.*** This is information that you can live without. Up to 90 percent of all information stored in file cabinets and databases is never retrieved, so this category can include a lot of data.

DETERMINE TYPE (OR TYPES) OF BACKUPS

Different types of data and different time-to-data requirements will require different backup processes and media. Types of backups include:

- Regular backup to tape or other removable media
- Remote mirroring
- "Electronic vault" storage via a Wide Area Network (WAN) or the Internet
- Periodic or real-time backup to a cloud provider

You will probably use a combination of techniques, balancing time-to-data versus cost trade-offs. Traditional tape backups are still widely used and can be effective, but they can create transportation issues, storage issues, and restoration issues. If not handled and stored properly, tapes can fail without warning. They require that the application also be reloaded, and software to read the tapes must be available. Electronic vault storage allows you to save your data over a WAN, such as the Internet, and can be easier to restore than tape. Remote mirroring ensures that there is little or no data loss, but it is the most expensive option. Cloud service providers can provide quick access to virtual servers and data if you are prepared in advance to quickly access them.

DEVELOP RECOVERY PROCESSES

Now you must develop and document the process for both backup and recovery of data. It does no good to have a plan in your head or one that sits on the shelf. Schedules will need to be developed to ensure that backups are made in a timely fashion. Some criteria to be considered when evaluating which recovery techniques to use include:

- ***RTO (Recovery Time Objective).*** How quickly must the data be restored before business is adversely affected?
- ***RPO (Recovery Point Objective).*** How much data can you afford to lose before the business is adversely affected?
- ***Availability.*** Can the system be down while you create the backups?
- ***Restoration.*** How sure do you have to be that you can restore the data?
- ***Value.*** How much is it worth to protect the data?
- ***Performance.*** What are the performance requirements of the application?

You must also consider how effective each recovery technique is in protecting from the different types of loss. Each business process may have a different recovery process.

DATA STORAGE OPTIONS

There are numerous options for data storage, each with its own advantages and disadvantages.

TAPE BACKUP

Tape backup is almost as old as computing itself. Tape has a low cost per gigabyte, and it is relatively easy to transport and store. Tape has been a reliable workhorse for the storage and archiving of important data, but it is not foolproof. Tapes can fail, so it is critical that backup tapes are periodically audited. The audit should be done by randomly selecting a tape and verifying that it can be read and restored using different equipment than that used to create it. An emergency is not a good time to discover that the tapes are unreadable or can only be read by the equipment used to create the backup.

If the data files are important enough to back up, then they are important enough for you to implement the appropriate levels of physical and logical security. Ensure that the tapes are stored in a climate-controlled location free of dust and other sources of contamination. You should also make multiple copies of the tapes that can be stored in different locations to increase the chances of the data surviving a disaster.

Almost as important as how and where the tapes are stored is creating a tape rotation schedule. It is impractical in all but the smallest organizations to back up everything each time a backup is performed, so the normal practice is to perform a full backup periodically (e.g., weekly) followed by regular backups of any changes that have occurred since the full backup. The most common tape rotation strategy is called the Grandfather-Father-Son (GFS) backup scheme. It offers the following benefits:

- A minimum number of tapes is required to keep all the system's data backed up.
- It is easy to understand and perform, making it more likely to be followed.
- It is relatively easy to restore lost data from backups using this process.
- It minimizes the wear and tear on both the tapes and the equipment.

The most common GFS backup process is to use a seven-day schedule where a full backup is created once a week (usually over the weekend). Incremental backups are then made each of the other days. Tapes can be rotated and reused at specified intervals, depending on how many copies you wish to store. An example GFS backup strategy is as follows:

Create an initial complete backup. Label this tape "Month 1" and store off-site. This is the first "Grandfather" tape.

Create a full backup at the beginning of the week. Label this tape 'Week 1" and store off-site. This is a "Father" tape.

On each of the other days, perform an incremental backup using a different tape for each day. Label each tape with the appropriate day of the week. These are the "Son" tapes.

On the same day of the week that you did the first full backup, perform another full backup and label the tape "Week 2."

Repeat for each week of the month, reusing the incremental backup tapes each week.

After four weeks, make a full backup and store this tape off-site. This becomes the second "Grandfather" tape. The first "Grandfather" tape can now be reused.

Repeat the weekly process, reusing the "Father" tapes from the previous month.

DISK MIRRORING

By writing data to two different disks you create two identical copies, which increases the odds that at least one copy of the data is available at all times. The main disk used to store the data is called the protected disk, and the disk to which the data are replicated is called the backup disk. The two disks can be in the same location or in different locations. A WAN is used if the backup disk is at a different location from the protected disk. Installing the backup at a different location provides protection against a disaster that occurs at the location of the protected disk. While disk mirroring is an effective approach, beware of its impact on your network traffic load.

Two different types of disk mirroring are available, synchronous and asynchronous. Each provides a different time-to-data recovery, and each has different performance considerations.

Synchronous mirroring works by writing to the backup disk first, then writing to the protected disk once it has been confirmed that the write to the backup disk was successful (see Figure 20-5). This type of mirroring ensures that the backup data are always up to date, but it is slower and more expensive than asynchronous mirroring. Special disk controllers are required to enable the two-way communication between the disks, and there is inherent latency between the writing of the data to the backup disk and waiting for the confirmation.

Asynchronous mirroring (or shadowing) works by sending the data to both the protected and backup disks at the same time (see Figure 20-6). It is cheaper than synchronous backup, and more than one system can write to the backup disk. It is also quicker, since the application does not have to wait for a confirmation on the write to the backup disk. The downside to asynchronous mirroring is that you cannot be guaranteed that the last transaction before a disaster was successfully written to the backup machine.

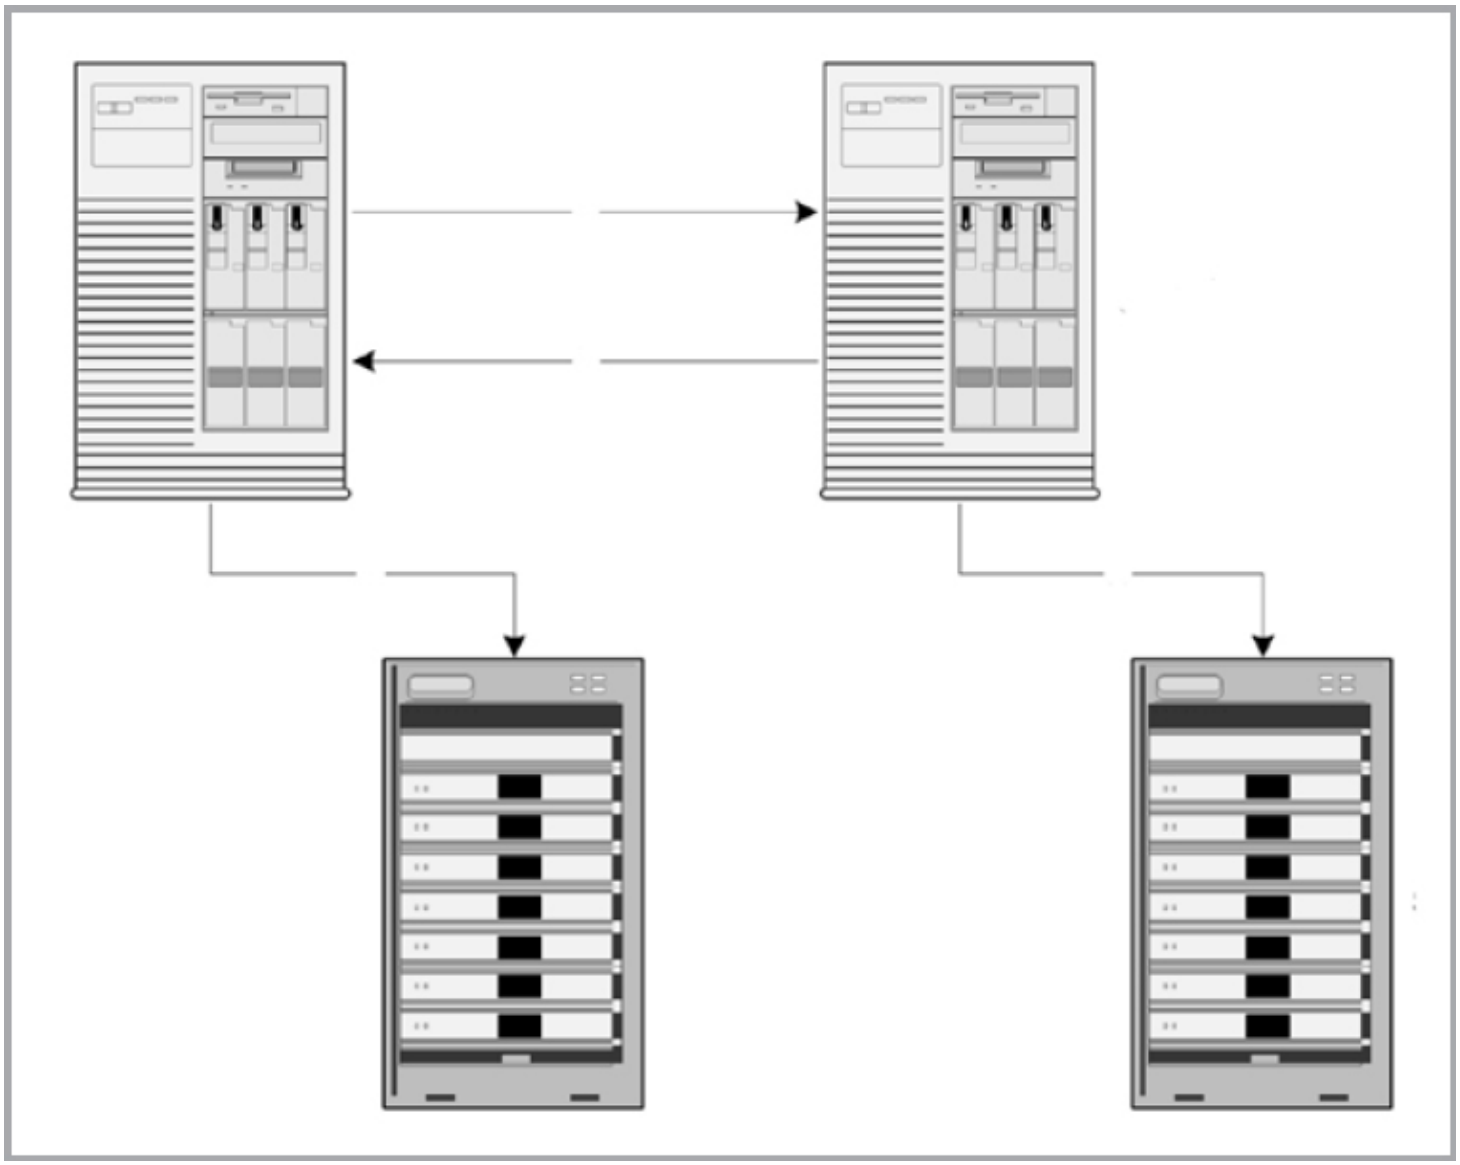

FIGURE 20-5. *Synchronous mirroring.*

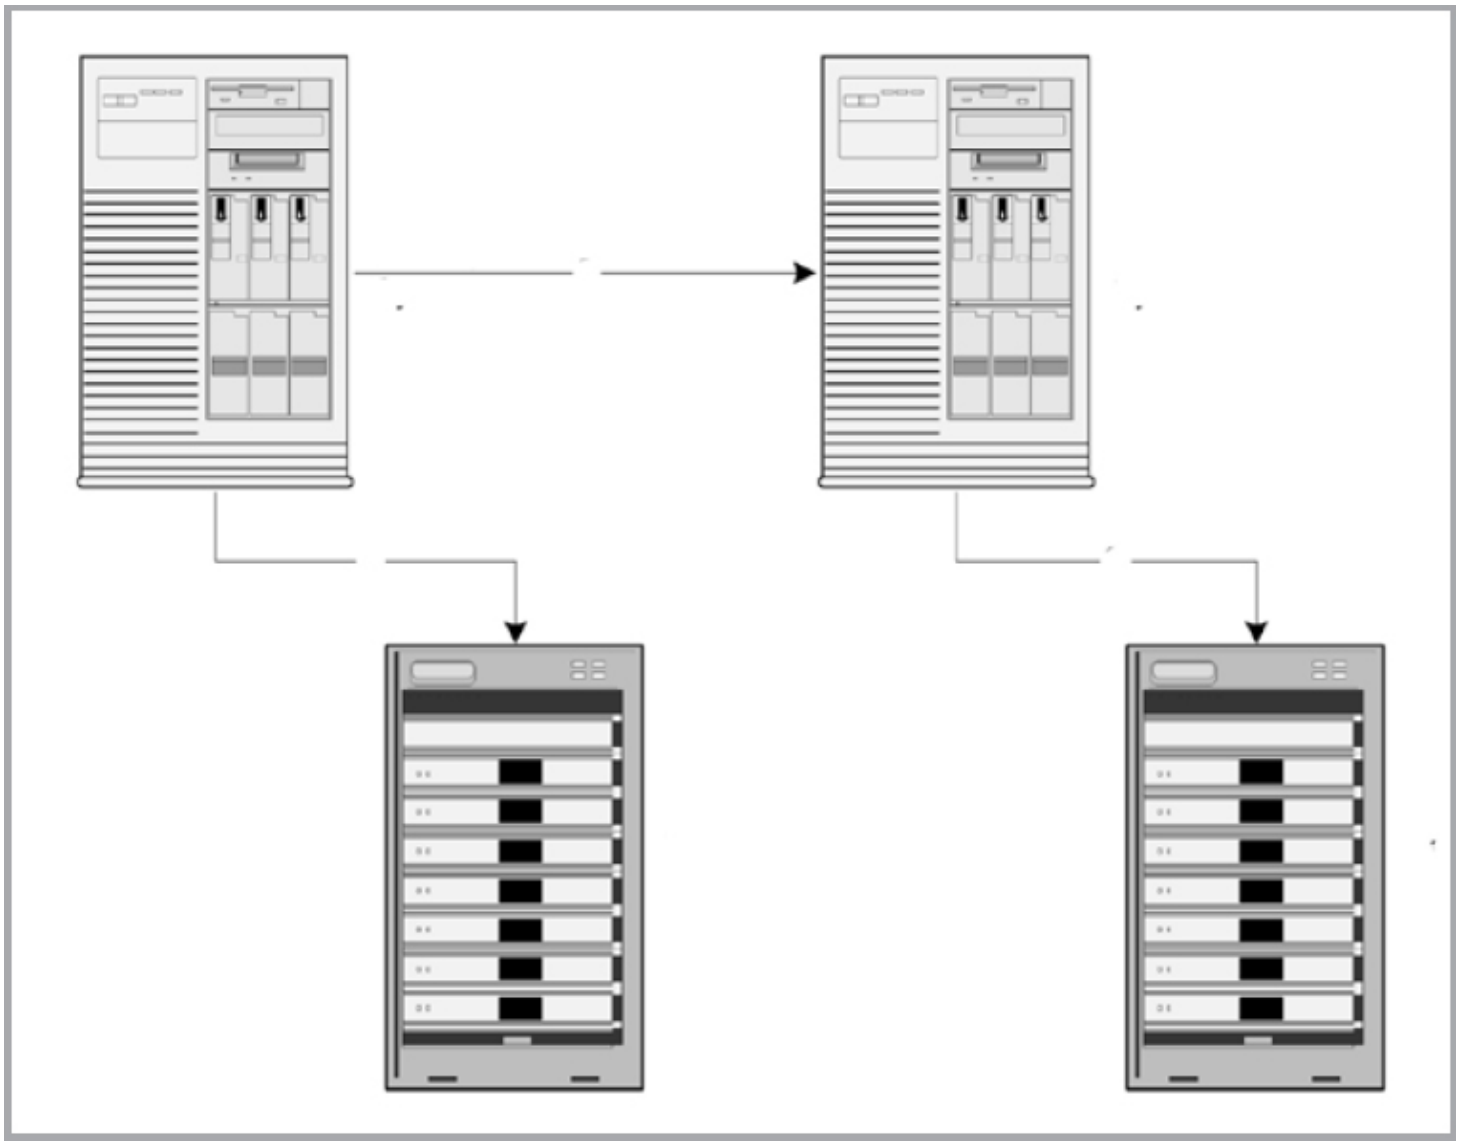

FIGURE 20-6. *Asynchronous mirroring.*

RAID

RAID is an acronym for redundant array of inexpensive (or independent) disks and is used to provide fault tolerance to disk storage systems. RAID works by combining a collection of disks into a logical array of disks using a special disk controller that does not require all disks to be functioning to maintain data integrity. It can be implemented using either hardware or software. The RAID drives are seen as a single device by the operating system. RAID also increases disk performance and reliability by spreading the data storage across multiple drives, rather than a single disk. The terms used when describing a RAID implementation are defined as follows:

- ***Duplexing.*** Disk duplexing involves the use of two RAID controllers writing the same data to two separate disks simultaneously. A system using duplexing can survive the failure of either a disk controller or a hard disk.
- ***Mirroring.*** Disk mirroring involves the use of a single RAID controller writing the same data to two separate disks simultaneously. A system using mirroring can survive the failure of either hard disk. Both duplexing and mirroring can slow down system performance because the data is being written twice.
- ***Striping.*** Striping involves breaking up the data into smaller pieces and writing the different pieces to multiple disks. The data may be broken up into bits, bytes, or blocks depending on the RAID implementation used. Striping is faster than either duplexing or mirroring.
- ***Parity.*** Parity is a way to achieve data redundancy without the disk space overhead of mirroring by storing logical information about the data being written to facilitate recovery. Parity is used with striping and requires at least three disks. The parity information is either stored across multiple disks or on a separate disk.

There are several levels of RAID operation, each with its own balance of redundancy, fault tolerance, cost, and complexity.

- **RAID 0.** Disk striping. This implementation of RAID divides the data among several disks, which allows for good performance, but with no redundancy. There is no protection against data loss if a disk were to fail. RAID 0 is not recommended for data recovery purposes.
- **RAID 1.** Mirroring and duplexing. This level of RAID involves mirroring or disk duplexing of the data across two or more disks, which provides for redundancy in case of a disk failure. Performance is slower than with RAID 0, especially during data writes. This level is simple and inexpensive to implement, but 50 percent of the storage space is lost because of the data duplication.
- **RAID 2.** Bit-by-bit striping. This level of RAID stripes data bit by bit across multiple drives and is used with disks without built-in error detection. Since most modern disks have built-in error detection, RAID 2 is rarely used today.

RAID 3. Byte-by-byte striping. This level stripes data byte by byte across multiple drives, with the parity information stored on a separate disk. The parity disk can be used to restore data if a failure occurs. The parity information is at risk because it is stored on a single drive.

RAID 4. Block-by block striping. This level of RAID stripes data at the block level. Just like RAID 3, the parity information is stored on a separate disk. Performance is greater than with RAID 2 or 3 because the data is handled in block sizes.

RAID 5. Striping with distributed parity. This level is similar to RAID 4, except that the parity information is stored among the available disks. RAID 5 is a common implementation of RAID.

RAID 10. Mirrored striping. This level of RAID (sometimes called RAID 0+1) is a combination of RAID levels 0 and 1. Data is striped across multiple disks and also mirrored. It provides the best fault tolerance of all the RAID levels but is obviously the most expensive.

LOAD BALANCING

Load balancing is used to distribute network traffic dynamically across a group of servers running a common application to prevent any one server from becoming overwhelmed. Using load balancing, a group of servers appears as a single server to an application on the network. The load balancing process is part of the network operating system; the process monitors each server to determine the best path to route traffic on the network to increase performance and availability. Load balancing also allows the application to continue running even if one of the servers goes down. As long as at least one server is available, the application will continue running. Load balancing can be implemented on different servers at a single location or at different sites. If load balancing is implemented on servers at different sites, it can act as a method to allow access to applications in the event of an incident at one of the locations.

NETWORK ATTACHED STORAGE (NAS)

A NAS environment is a common storage area for multiple servers. NAS environments are useful for storage or file server applications, such as mail and web services. A NAS server runs a minimal operating system, and is optimized to facilitate the movement and storage of data. Using a NAS environment creates a centrally managed storage pool, which allows new storage devices to be added without requiring network downtime. Storage volumes from a down server can be easily reassigned, or new storage can be easily added if needed. The flexibility provided by a NAS environment increases the availability and reliability of network storage, adding value to your disaster contingency plans.

STORAGE AREA NETWORKS (SAN)

A SAN is a high-speed, high-performance network that allows computers running multiple operating systems to store data on a single virtual storage device. A SAN is designed to handle backup traffic more efficiently than a NAS environment. The SAN can be local or remote and usually communicates with the server using a fiber channel. By moving the storage off the Local Area Network (LAN), backups can be performed without affecting the performance of the applications on the LAN.

CLOUD BACKUPS

Online data storage is becoming popular as cloud services are being used more and more to provide services to end users. With a cloud provider, changes to your data are delivered via the Internet to the cloud provider. This allows your data to be stored at a secure, professionally managed location away from any dangers to your facility. Cloud servers and data storage offer the following benefits for disaster recovery:

- The most obvious benefit is that your data is automatically stored to another location. Backups do not have to be manually started or managed.
- Your data can be protected using a single solution that is accessible anywhere there is an Internet connection.
- Remote offices and road warriors can back up their data without requiring separate hardware or complex Virtual Private Network (VPN) solutions.
- There are lower upfront costs to implement an online backup solution—service is usually provided on a subscription basis.
- No special in-house technical skills are required as they become the cloud service provider's responsibility.

As with anything else, there are trade-offs you must be willing to make to implement a cloud data storage solution for use in disaster recovery:

- There are likely higher overall costs since the service is subscription based—the provider recoups its equipment and software costs over the term of the subscription agreement.
- There may be issues with retrieving your data from the provider.
- Your provider could experience an outage which prevents you from accessing your data.
- Restoring large amounts of data over the Internet consumes a lot of time and bandwidth.
- There's a risk that your cloud data storage vendor goes out of business. Then, what happens to your data?

A hybrid option for online backup combines the best of cloud backup storage with the best of traditional online backup services. Some vendors will provide you with an appliance and software that allow you to use the Internet to do online backups, yet still have physical access to the backup device at a location that you control. This protects you against some of the disadvantages of service-only online backup solutions, such as losing your data to bankruptcy of your service provider. It also makes complete restorations easier, as the device can be physically brought onsite and connected directly to the local network for quick restoration.

Figure 20-7 is a comparison of the relative availability value of the different data storage options described here.

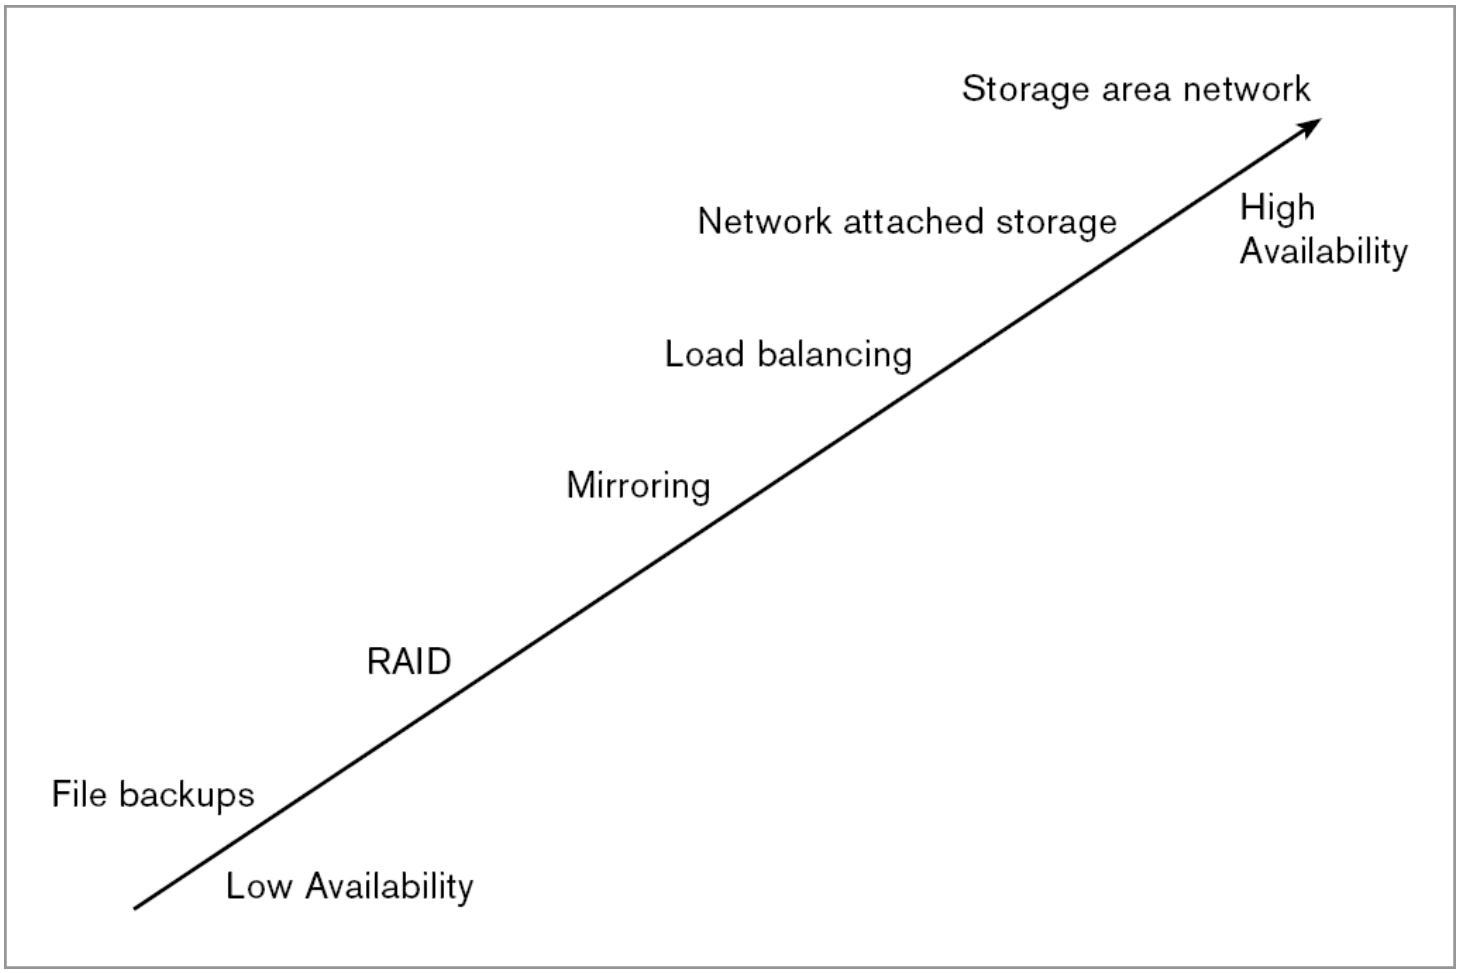

FIGURE 20-7. *Data availability options.*

VIRTUALIZATION

An increasingly popular technique for managing the explosion of data storage is the use of virtualization. Virtualization as it is mostly used today essentially relies on software to mimic hardware. The typical software application wastes a tremendous amount of storage space. Many applications require that some minimum amount of storage be allocated for their use, and then in many cases only a small fraction is actually used. Storage virtualization allows the physical storage from multiple storage devices to appear to be a single storage device to the operating system or application. The storage can then be allocated as needed for use by users, applications, and servers. Storage virtualization can increase utilization to 75 percent or better. (See Figure 20-8.) There are three basic approaches for virtualizing storage. The most common (called in-fabric) is the use of SAN devices connected using a high-speed fiber channel network. Software is then installed on a host server or a storage virtualization device is installed as part of the SAN; either provides the layer of abstraction between the hosts performing the I/O and the storage controllers providing the storage capacity.

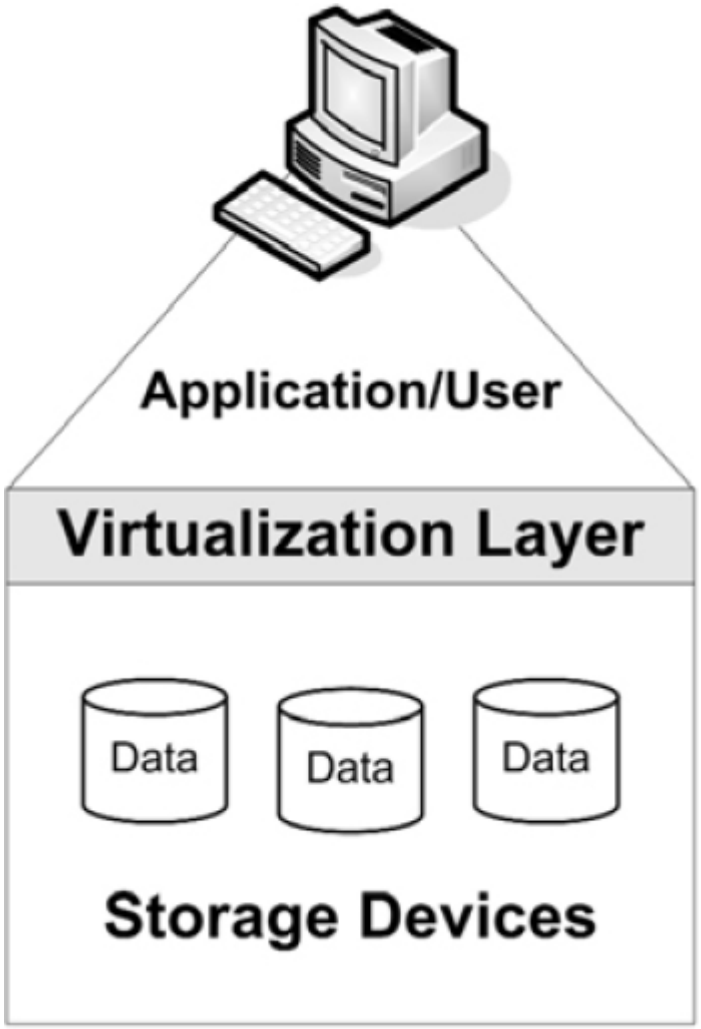

FIGURE 20-8. *Storage virtualization.*

The second method, called host-client, uses software running on file and application servers to detect available storage and to maintain the metadata necessary to manage them. The third method, known as in-array or embedded functionality, uses a special network controller and management software to manage different storage systems as one large resource.
All methods provide the advantages of storage virtualization, which from a disaster recovery prospective means:

- Data storage becomes more mobile and easier to deploy.
- Virtual tape library (VTL) technology can be used to decrease total backup time.
- Production and recovery storage environments no longer need to be strictly homogeneous.
- Disaster recovery costs can be lower, as less expensive storage devices can be used at the recovery site. The less expensive devices, while maybe not ideal, might work fine until the production devices have been restored.
- Administration is made easier during recovery as virtualized storage can be managed from a single administrative console.
- There is greater flexibility in managing changing application storage requirements. Issues to be aware of when considering storage virtualization include:
- If you have heterogeneous storage devices, are they all compatible with the virtualization technology you are considering?
- Applications could experience a decrease in performance if the recovery site data storage hardware has a slower response time.
- Measures must be taken to ensure that the storage metadata is protected and backed up.
- Be aware that whichever option you choose, you'll be locked into a particular vendor.

PLAN TESTING AND MAINTENANCE

Your plan will need to be tested and updated periodically to keep it effective. Testing allows you to identify any deficiencies in the plan so that they can be corrected. This includes not only equipment and backup issues but personnel issues as well. Each component of your plan must be tested to verify the accuracy and completeness of your recovery procedures and to determine the overall effectiveness of the plan. Some areas to review when testing your plan include:

- Ability to restore critical applications from backups
- Performance of recovery personnel
- Performance of backup equipment

- Communications

To remain effective, your plan must be kept up to date with changes in your production systems. Changes in your IT infrastructure can be caused by business process changes, technology upgrades, regulatory requirements, employee turnover, or new policies. Any of these events should trigger a review of your recovery plan. It is therefore essential that your plan be reviewed and updated frequently to ensure that new processes are documented and that the recovery process is updated to reflect these changes. Items to be monitored for changes include:

- Hardware, software, and peripheral equipment
- Business operation requirements
- Security requirements
- Technology changes
- Recovery team contact information
- Vendor information
- Regulatory requirements

Because your recovery plan contains potentially sensitive personnel information, its distribution must be controlled. Copies should be stored at the home of key recovery personnel, at your production location, and at your off-site recovery location with your backup media. The individual in charge of recovery planning must maintain a list of who has copies of the plan and a record of when changes were made. (Use Form 20-1, Recovery Plan Distribution List, and Form 20-2, Recovery Plan Change Record, found on the companion url, as examples.)

CONCLUSION

Data is the lifeblood of modern businesses; by having an effective data recovery plan you can help ensure that your business will survive an unexpected emergency. The steps are simple, but must be diligently performed to be effective:

- Identify what data is important.
- How soon do you need it?
- What is it going to cost not to have it?
- Test your recovery procedures