Realtime publishers

The Shortcut Guide™ To

Availability, Continuity, and Disaster Recovery

sponsored by

ARCserve.com™

Dan Sullivan

Chapter 3: Top-5 Operational Challenges in Recovery Management and How to Solve Them .. 33
Challenge 1: Scheduling and Monitoring ... 33
Scheduling: Various Types of Backups and Their Uses ... 34
Types of Backups ... 34
Backup Scheduling ... 36
Monitoring ... 36
Monitoring Backup Operations ... 37
Verifying Backups ... 37
Challenge 2: Choosing the Right Storage Media Options ... 38
Tapes: Advantages and Disadvantages ... 39
Disks: Advantages and Disadvantages ... 39
Identifying the Best Option for Your Business ... 40
Step 1: Classifying Data by Protection Requirements ... 40
Step 2: Mapping Recovery Objectives to Backup Options ... 41
Step 3: Determine the Most Cost-Effective Ways to Meet Recovery Objectives ... 41
Challenge 3: Controlling the Costs of Offsite Storage ... 42
Time to Move Data to Offsite Storage ... 42
Risks Associated with Offsite Tape Rotation ... 42
Cost of Offsite Storage ... 42
Cloud Storage as an Offsite Storage Option ... 43
Challenge 4: Keeping Up with Growing Volumes of Data ... 43
Deduplication Options ... 44
Challenge 5: Recovering When Disaster Strikes: Continuity and Failover ... 45
Physical Requirements ... 46
Application Requirements ... 47
Disaster Recovery Service Providers ... 47
Summary ... 47

Copyright Statement

Chapter 3: Top-5 Operational Challenges in Recovery Management and How to Solve Them

Maintaining effective recovery management procedures is not a trivial task. From making sure processes are running correctly to controlling the costs of operations, there is no shortage of challenges. In this chapter of the *Shortcut Guide to Availability, Continuity, and Disaster Recovery,* we will examine five of the top operational challenges we commonly face in recovery management:

- Scheduling and monitoring
- Choosing the right storage media option
- Controlling the costs of off-site storage
- Keeping up with growing data volumes
- Recovering when disaster strikes

These five challenges are interrelated. For example, dealing with growing data volumes is directly related to controlling costs. Monitoring is less directly related to recovery operations but is just as important—we do not want to find out about a failed backup operation when we try to restore critical data after a disaster. As we address each of these five challenges, we will consider both the fundamentals of the individual challenges as well as how the challenges influence each other.

Challenge 1: Scheduling and Monitoring

Scheduling and monitoring go hand in hand. It is the combination of deciding what types of backups are required and making sure they are performed correctly. As with most IT operations, recovery management is driven by business requirements, so in the course of this discussion, we will have to refer back to those requirements when deciding on the proper schedule of backup types.

Scheduling: Various Types of Backups and Their Uses

In Chapter 1, we looked at the business case for management recovery. That discussion included some obvious requirements, such as restoring from isolated failures and compliance, and some not-so-obvious requirements, including varying backup policies based on the business value of data. Meeting these requirements with operational procedures entails implementing the right combination of backup types at the right times.

Let's start with a quick review of the various types of backups, then consider how we combine them to achieve the level of data protection we need.

Types of Backups

There are a few different types of backups because we need to balance competing requirements in recovery management. Ideally, we would have comprehensive coverage of all data at all points in time and we would be able to restore any or all of that data rapidly. Add to that minimizing cost and resources required to perform backups, and we would have the ideal solution. We will not be getting our ideal solution anytime soon, so we will have to settle for the optimal realizable solution.

Pragmatic recovery management solutions build on three types of backups to find that optimal solution:

- Full backups
- Incremental backups
- Differential backups

These methods have different advantages and disadvantages and so tend to complement each other as part of a recovery management strategy.

Full Backups

Full backups, as the name implies, make a complete copy of data that might later need to be restored. It has a number of advantages. A full backup is a completely self-contained backup, so a complete restore operation can be performed using a single backup. A key disadvantage of full backups is the time and space required to create and store them. The size of full backups is proportional to the amount of data to be protected. (The size is proportional to, not equal to, the size of the source data because of compression).

Incremental Backups

Incremental backups reduce the time and space required for full backups by copying only data that has changed since the previous backup. Consider a simplified example to see how significant the storage reduction can be.

If we assume that only 10% of data changes, then the amount of storage required for backup can be significantly reduced. For example, assume we have to backup about 5TB of data. Assuming a 30% compression rate by the backup software and a 10% rate of changes, an incremental backup of 5TB of data can be stored with as little as 350GB of storage (see Table 3.1).

Total Data Size (GB)	Full Backup Size (GB) (30% Compression)	Changed Data Size (GB) (10% Rate of Change)	Incremental Backup Size (GB) (30% Compression)
500	350	50	35
1000	700	100	70
2000	1400	200	140
5000	3500	500	350

Table 3.1: Incremental backups yield significant savings in storage when compared with full backups.

Clearly the savings in storage is substantial; however, in IT as in economics, there are no free lunches. What we gain in reduced storage costs and time to perform backups is accompanied by a disadvantage during restore operations.

In its simplest form, restoring from incremental backups requires restoring from a full backup and then restoring each incremental backup performed since the last full backup. Depending on backup software, it is possible to create a synthetic backup, which is a backup that results from merging a full backup and some number of incremental backups. The advantage of synthetic backups is that they combine the advantages of both full and incremental backups. As Figure 3.1 depicts, a synthetic backup is a full backup plus all changes since the backup was made.

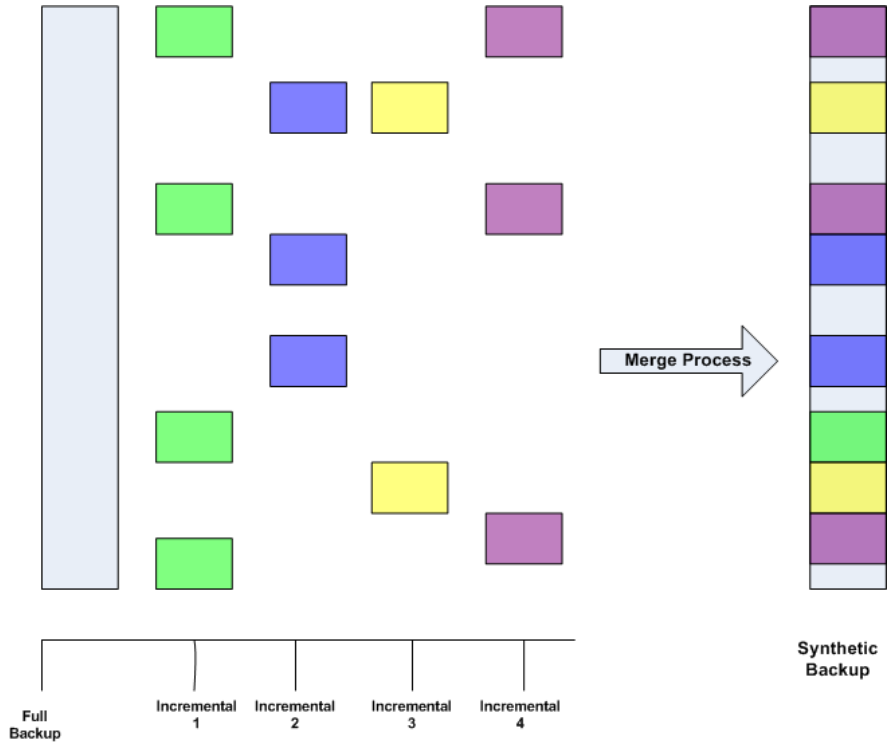

Figure 3.1: Synthetic backups merge full backups and incremental backups to allow for more efficient restore operations.

Differential Backups

An alternative method to incremental backups is the differential backup. Like incremental backups, differential backups start with a full backup and then back up only changes. The difference is that the differential backup captures all changes since the last full backup, not since the last backup whether it was full or incremental. An advantage of differential backup is that it only requires the full backup plus the latest differential backup to perform a complete restore. A disadvantage, relative to incremental backups, is that additional storage is needed for the differential backup.

Given the various types of backups, what is the best combination of these backup methods to provide adequate data protection while controlling costs?

Backup Scheduling

There a number of ways to schedule backups that are based on the idea of a rotation. The idea started when tapes were virtually the only option for most businesses, so we sometimes hear these methods referred to as "tape rotation" schemes. Rotation schemes balance the desire to re-use backup storage while maintaining a reasonable number of recovery points.

One of the simplest schemes is the round-robin method. A business has a fixed amount of backup storage, either tapes or disk, but we will describe the process in terms of tapes for simplicity. Under the round-robin scheme, a tape is used for each backup and once all tapes are used, the tape containing the earliest backup is re-used. A small business, for example, may be satisfied with using five tapes and performing a full backup each night of the business week. Every Monday, the previous Monday's tape is overwritten, and so on for each day. This method is simple to manage but leaves the business without a recovery point earlier than the prior week. This method makes it difficult to maintain off-site backups while preserving the ability to rapidly recover in case of an isolated failure.

A common rotation scheme is based on a combination of monthly, weekly, and daily backups. This is sometimes referred to as the grandfather-father-son (GFS) method. The basic idea is that a full backup is done each week (the father) and differential backups are performed each day (the son). At some point in the month, the latest father tape is promoted to "grandfather" status and removed from the rotation. This backup can be stored for long periods in off-site storage. The remaining weekly backups are maintained, typically on-site, and overwritten after 1 month. The daily backups are overwritten on a weekly basis. There are a number of variations on the GFS scheme, but the key points are that it supports long-term off-site backups as well as weekly and daily backups.

The GFS scheme provides a good balance of efficient use of tapes or disk with the ability to maintain multiple recovery points.

Monitoring

Monitoring backups is a two-part process. We want to monitor backup operations to ensure they execute as expected and we want to verify that the backups performed are valid and usable.

Monitoring Backup Operations

As the number of servers that requires backup grows, so does the complexity of monitoring and managing backup operations. The complexity stems from a number of factors:

- Not all data is equally valuable to the business. Different types of data require different levels of protection, and that means different backup schedules and retention policies.
- As the number of servers and data volumes grows, so does the likelihood of a failure somewhere in the overall operation. For example, if a failure is likely to occur 1 in 1000 backup operations, and we run 10 backup operations at a time, the chances of having a failure somewhere in the enterprise is 1 in 100.
- Data is geographically dispersed. Small remote offices will require data protection services but will likely not have a dedicated IT staff on-site placing additional demands on central IT professionals.
- There are simply more devices and processes to manage as data volumes and the number of servers grows.

As a baseline, backup administers should capture and monitor several metrics for all backup operations:

- The start and end time of backup operations—Duration information can help identify trends in changes to the amount of time required to perform the backup.
- The amount of storage used for backups—As with time, this can be used to detect long-term trends, but it can also help optimize storage use in the short term. For example, backups with modest storage requirements may be combined to make more efficient use of storage devices.
- Performance metrics on backup servers, particularly memory use and CPU utilization—Although backup and restore operations are obviously I/O intensive, compression and encryption can put significant load on the CPUs of the backup server.

In addition to those metrics, it can be important to track errors and failures. Significant errors, such as a failed backup, should be addressed as soon as possible in most situations. Minor failures, such as the detection of a bad block on a disk, should be monitored over time in case the failure is not an isolated incident but part of a larger systemic problem.

Verifying Backups

When you need to restore a file from a backup, it is the wrong time to find out the backup is corrupted. There are few ways to verify backups, although details will vary with backup software.

One way to verify backup is to have the backup software perform a verification pass immediately following the backup operation. This operation compares the contents of the original files with the files on the backup medium. The ability to perform this type of verification will depend on your backup software. It will also extend the time required to perform a backup, which may not be an option if you are already working within a tight window of time.

Compression features may also provide the ability to check the integrity of the compressed files without performing a block-by-block comparison. This could be faster than other full comparison operations but again depends on the features of your backup software.

A tried-and-true method is to randomly select a backup, restore it, and compare it with the original file. This is helpful but does have some drawbacks. First, unless a backup administrator is well versed in statistics and probability, she might not adequately sample the set of backups at least from a formal quality control perspective. Second, data is constantly changing, so the original files may not be available to compare against. (After all, one of the reasons we make backups is to be able to restore a previous state). Even with these limitations, it is useful to perform some level of verification using features offered by backup software and to randomly select backups for manual verification.

Scheduling the right combination of full, incremental, and differential backups helps to ensure a business will be able to recover to a number of points in the case of isolated or catastrophic failure. The verification process is an extra but worthwhile step to help ensure the integrity of backups.

Costs are always a factor that must be taken into consideration when managing risks, and backups are no different. An optimal backup schedule will balance costs with the ability to recover to different points in time. Verification processes may be less than ideal because of other business constraints, including costs. Nonetheless, proper scheduling and monitoring of backup operations contributes to data protection and overall risk mitigation related to data loss.

Challenge 2: Choosing the Right Storage Media Options

Magnetic tapes have a long history in information technology, and they have proven themselves a reliable and cost-effective storage media. As volumes of data have grown along with the seeming never-ending desire for faster backup and restore operations, the information technology industry adopted disk storage as an alternative. There are advantages and disadvantages to both, so choosing between them is not a trivial task. The key to finding the right choice, or the right combination of tape and disk, will depend on a sound understanding of operational requirements, budget constraints, and your business' data protection strategy.

Tapes: Advantages and Disadvantages

Tape backups have a number of features that make them a clear choice for business backups:

- Tapes are cost effective—The industry Linear Tape Open (LTO) standard is widely adopted, so businesses have a non-proprietary option with LTO, which often translates into competitive pricing. (LTO products often use the term Ultrium as well). The cost per gigabyte of storage is generally lower for tapes than for other media, including disks.
- Tapes are easy to physically transport—This is especially important with regards to archiving and offsite storage.
- Tapes have a long history as the backup medium of choice, so tapes are compatible with many backup applications.

Now, let's consider for the disadvantages:

- Tapes use a sequential access method, which results in relatively slow restore operations when compared with disk-based restores.
- Tapes are more vulnerable to environmental factors, such as heat and humidity, than disks. Environmental factors can adversely affect the durability and reliability of tapes.
- Tapes are routinely physically handled, which opens opportunities for human error, such as lost tapes, and unintentional damage.

The fact that tapes are easy to handle is an advantage as well as a disadvantage shows the kinds of challenges and lack of clear-cut choices we have with backup media.

Often, the advantages of tapes outweigh their drawbacks. Advances in disk technology created a viable alternative to tapes.

Disks: Advantages and Disadvantages

Disk storage has a number of advantages for data protection. Some of these, not surprisingly, correlate with disadvantages of tapes. The advantages of disks as a backup media include:

- Disk performance is faster than tape—This is due to both the random access nature of disks versus the sequential access pattern of tapes and to the speed with which data can be written to disks. Disk drives used for backups often use fibre channels, which have higher throughput than tape drives.
- Disks are more protected from environmental factors than tapes, which are typically moved from tape drives to short- and long-term storage locations.
- Disk capacity can grow relatively easily when using a storage area networks (SAN).

Disk storage does have its downsides:

- Disk are generally more expensive than tapes when comparing the cost of storing a gigabyte of data
- Disks may not be suitable for archiving, in which case, mobility of the storage media is especially important

So how should one balance the pros and cons of both options to come up with the optimal mix?

Identifying the Best Option for Your Business

The best option for your business is determined by balancing often-competing requirements. For example, maintaining a low-cost solution is an obvious concern, but the total cost of data protection includes more than the cost of tapes and a tape drive. If there are significant opportunity costs associated with long recovery times, backing up to disk may actually cost the business less in the long run than using tapes. However, backing up archived user directories to tape may provide the required data protection at a lower cost than using disks because rapid restoration is not essential.

Step 1: Classifying Data by Protection Requirements

The first step to identifying the best option for your business is to classify data according to the level of protection required for each type of business data. In this case, the level of protection includes:

- Recovery point objectives (RPOs)
- Recovery time objectives (RTOs)
- Archiving requirements
- Other compliance-related requirements

Some data will require near-time recovery points and rapid recovery, such as an order management system. For these data, the additional cost of disks can be justified. For other types of data, the reliability of the storage media may be especially important. Although email may appear at first glance to be a moderate concern, that can easily change in the event of litigation. If e-discovery is required in response to litigation, a business should be able to produce required documents from online systems or archives.

> **Resource**
>
> For more information about e-discovery, see the Electronic Discovery Reference Model at http://edrm.net/.

Also, consider the need for data life cycle management. The amount of data stored cannot grow forever without cost; at some point, the value of retaining archives will be outweighed by the cost of maintaining those archives. The difficult process is determining when one reaches that point.

Step 2: Mapping Recovery Objectives to Backup Options

For each category of data, determine the volume of data that must be processed in a given time period to meet RTOs. Also, determine the amount of storage required to meet RPOs. These two pieces of information are somewhat dependent on each other. For example, the amount of data that needs to be processed to meet recovery point and time objectives will depend on which type of backup is used. You might want to consider a few different scenarios with different combinations of full, incremental, and differential backups to understand the trade-offs between time to backup, time to restore, and the volume of storage needed for backups.

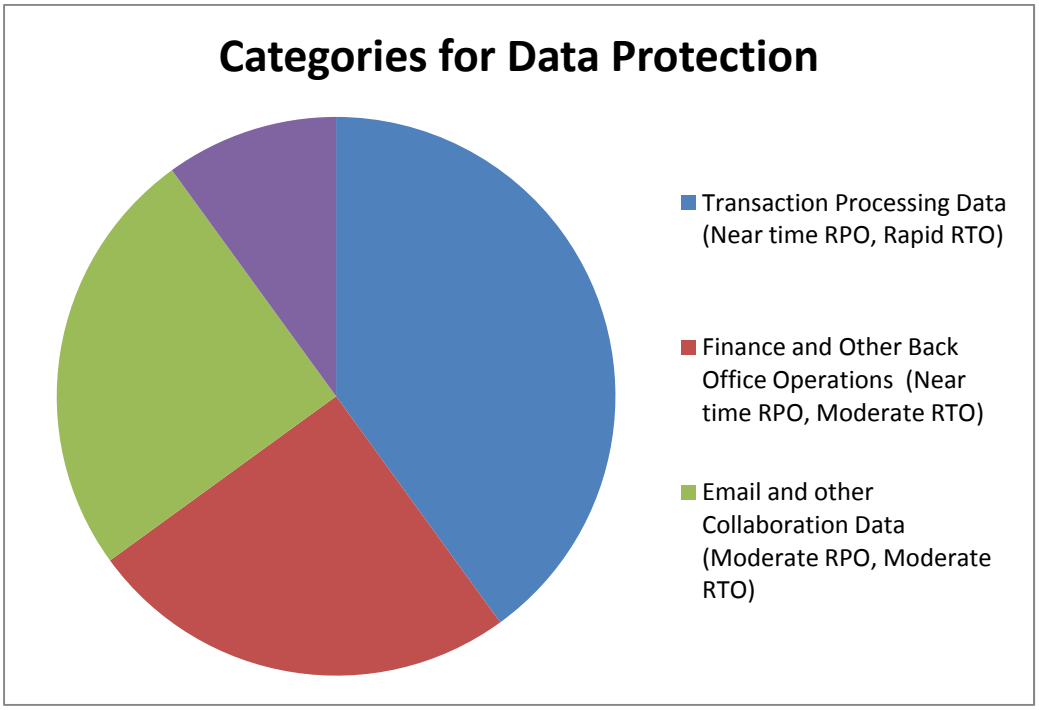

Figure 3.2: Data should be categorized according to recovery and archiving objectives to facilitate determining the optimal use of disk and tape storage.

Step 3: Determine the Most Cost-Effective Ways to Meet Recovery Objectives

The last step in the process is selecting the most cost-effective storage option that meets your requirements. If an option does not meet the needed recovery point or time objectives, the question of cost is irrelevant.

Be sure to consider a combination of tape and disk options to meet all requirements. For example, it may be best to use disk-to-disk backup to meet recovery objectives and then use tapes for archiving. When the archives are created from the backups, the process is called "disk to disk to tape." One of the advantages of this approach is that the time to backup up from a production system is reduced, but you still retain the long-term cost advantage of using tapes. In addition to cost advantages, this method can reduce the load on servers being backed up by allowing archives to be made from a copy of production data. Another cost consideration with regards to archives and disaster recovery is the cost of offsite storage.

Challenge 3: Controlling the Costs of Offsite Storage

Offsite storage is an essential part of comprehensive data protection strategies. It provides a level of protection in the event of a disaster at a primary facility. Businesses will vary in their use of offsite storage and the types of facilities used. For example, for a small business, offsite could be a safe deposit box in a local bank. Larger businesses might use different offices to store each other's backup tapes or they may use a third-party service provider. Regardless of what form of offsite storage is used, there are a number of considerations to controlling the cost associated with this part of data protection.

Time to Move Data to Offsite Storage

The time required to move data to offsite storage, and the associated cost, will depend on the type of storage media. Moving data to offsite disk storage, such as a backup service that hosts disk arrays for clients, will have relatively low labor costs. Depending on the volume of data transmitted, there may, however, be a marginal cost for necessary network bandwidth. The time required to move the data will depend on the volume of data and the network speed. Here again, RPOs and RTOs will help determine the necessary capacity.

The cost of moving tapes to offsite storage will be typically dominated by labor costs. Small businesses may be able to use informal practices, such as having an employee leave the office early to drop off tapes at another office on the way home at the end of the day. Other businesses will require more established procedures using staff or third-party providers to transport tapes.

Risks Associated with Offsite Tape Rotation

Any time tapes are moved offsite, there is a risk of losing or damaging the tapes. A lost tape could be a minor inconvenience or a significant problem, depending what is on the tape. If the tape contains sensitive information, such as personally identifying information, protected health care data, or financial information, it should be encrypted to prevent unauthorized disclosure. Consider compliance requirements when transporting tapes offsite to ensure information is adequately protected when it leaves the business.

Cost of Offsite Storage

The cost of offsite storage will vary with the volume of tapes stored offsite, the length of time they are stored, and possibly the number of times the storage facility is accessed. In addition, the quality of the storage facility with regard to data protection will influence cost. Renting a self-service storage locker from a local vendor might be the least-expensive option possible, but we can forget about any type of environmental controls, fire suppression equipment, or other necessary controls. At the other end of the spectrum, storing tapes inside a mountain may protect your tapes as well as one can expect, but the costs for such a service may only be justified for the most sensitive corporate information.

Cloud Storage as an Offsite Storage Option

Cloud storage is an emerging option for offsite storage. Public cloud providers typically charge by the volume of data stored and the length of time the data is stored. This type of "pay as you go" model may help reduce the need for additional hardware as the volume of data grows. As with other third-party providers, we must consider the long-term viability of the cloud storage provider along with the reliability of its service. Also, consider using strong encryption when storing confidential, sensitive, and private data in the cloud. Remember that the definition of strong encryption changes with time. Encryption algorithms and key lengths that were once consider sufficient for protecting information can be compromised using today's hardware and cryptanalysis techniques. The cost of offsite storage depends on a number of factors, including the volume of data we store, what type of media is used, the level of service provided by the storage facility, and how the data is transported to the facility.

Challenge 4: Keeping Up with Growing Volumes of Data

Businesses are faced with growing volumes of data. When we look into the source of this phenomenon, we find there is no single source of additional data; instead, it is driven by a wide range of business initiatives and requirements:

- Additional applications designed to take advantage of new opportunities that generate additional data
- Increasing use of collaboration tools, including email, messaging, and document repositories
- Compliance requirements that specify data retention requirements
- Improved analytics that drive the motivation to collect more data in order to more efficiently and effectively target key markets and customer segments
- The ease with which we can all download documents, presentations, videos, and other business-related content from the Internet

Much of this data will fall under the umbrella of enterprise data protection strategies, which means more data to back up. It does not necessarily mean more time to back it up or longer recovery times or even more budget to cover the cost of additional hardware and storage media. About the only option left is to improve backup technology. A significant advance in this arena has been in the area of data deduplication.

Deduplication Options

Deduplication takes advantage of the fact that the contents of data blocks on disk are often duplicated across multiple data blocks. Duplication of data creeps into even the most storage-conscious organizations. In part, this is due to the fact that it is often easier to produce and manage information if we duplicate information. Consider a few examples: Multiple employees save the same copy of a report. Software developers save a number of versions of similar application code. Sales reports with common corporate and department data are sent to sales staff along with their personalized reports. This type of duplication creates opportunities for more efficient storage on backup media.

Deduplication works at the data block level. During a backup operation, a data block is read and compared with all other data blocks that have been read before it in the same backup operation. This step is faster than it may sound at first. Instead of explicitly comparing blocks with each other, a value (known as the value of a hash function) is calculated for each block. If a prior block has the same value, the backup program only needs to store a pointer to the previous block instead of storing another copy of the block. This can result in significant savings in storage.

There are two of options when performing the deduplication process, both of which have advantages and disadvantages:

- Source-side deduplication, in which the deduplication process runs on the server hosting the data being backed up
- Target-side deduplication, in which the deduplication process runs on the backup server

With source-side deduplication, the network traffic is reduced because duplicate blocks are not sent to the backup servers—only references to an already copied block are sent. A potential drawback is that the source server's performance is adversely affected by the additional load. This can be mitigated by scheduling backups during non-peak demand periods. However, if the source server is a virtual machine, you should consider the load on the physical server by other virtual machines during backup. When backing up virtual servers, be sure to consider the particular needs of virtual environments.

> **Cross Reference**
>
> See Chapter 2 for more information on this topic.

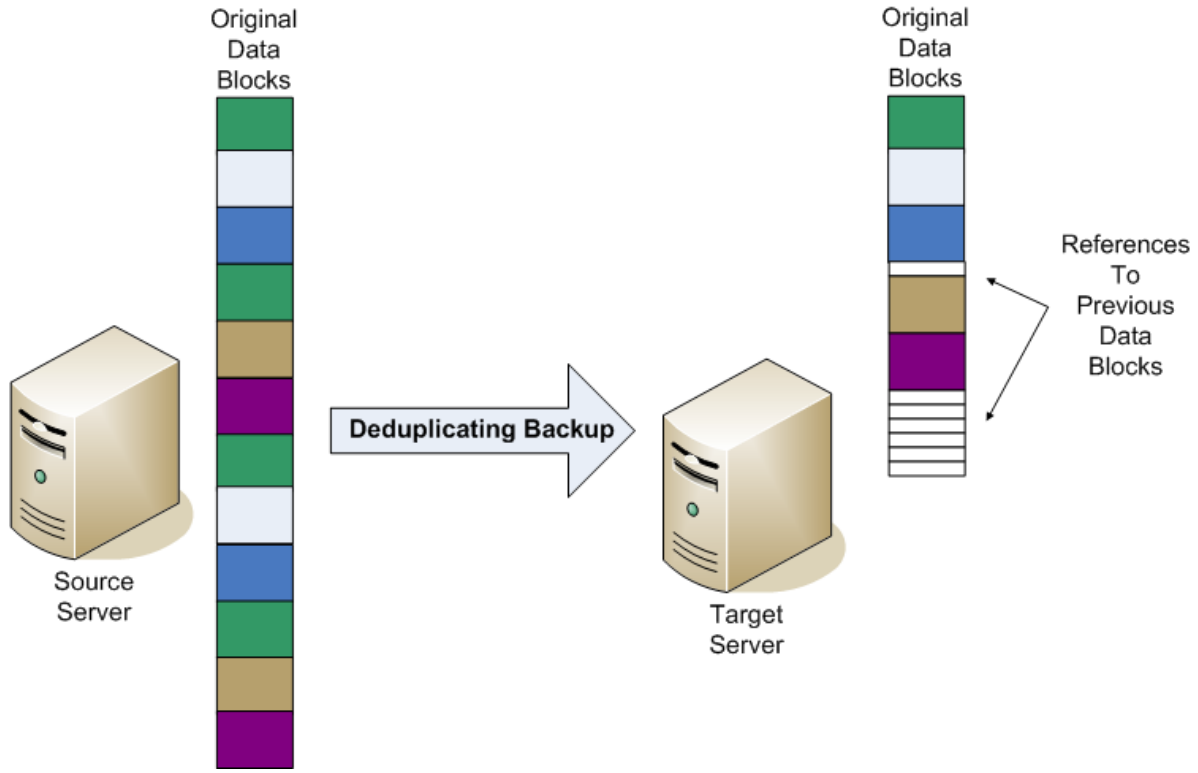

Figure 3.3: Deduplication reduces the number of data blocks stored in backups by substituting pointers to previously backed up data blocks.

Target-side deduplication performs the deduplication operations on the backup server. This approach can offer greater throughput in highly distributed environments with dedicated backup servers. A tradeoff is that network traffic is not reduced because duplicated blocks are copied from the source system to the backup server.

Deduplication is a key technique for accommodating growing data volumes in data protection strategies. With deduplication, existing backup infrastructure can keep up with growing data volumes without requiring an equivalent increase in storage media.

Challenge 5: Recovering When Disaster Strikes: Continuity and Failover

The last of the five key operational challenges of recovery management we will examine is disaster recovery. A business' disaster recovery plan answers the question "How would we keep the business running if there was a catastrophic loss at a key business site?" It is important to understand how this differs from other reasons we perform backups. If someone accidently deletes an important file, someone in tech support can restore the file from a recent backup. The staff is in place, the backups are available, the backup applications and servers are up and running, and, perhaps most importantly, it is an isolated incident. It may be an important file, but chances are its loss would not adversely affect ongoing operations. Disaster scenarios are different.

Disaster recovery plans are implemented when normal business operations cannot continue as normal. This can be caused by a range of factors:

- A catastrophic physical loss, such as a fire or natural disaster that destroys buildings housing the business
- An extended loss of power, including exhausting backup power supplies
- Regional storms that prevent employees from reaching offices and data centers

In cases such as these, normal backup and restore procedures are not enough and the chances of maintaining normal business operations will depend heavily on how well the business planned prior to the disaster. That planning should take into account both physical requirements and application requirements.

Physical Requirements

The physical requirements for disaster recovery include both information technology infrastructure and a physical space to house equipment and employees. At minimum, this includes:

- Backup servers for running essential applications—Business processes should be prioritized to identify which services should be restored first and which can wait. Also, consider the level of performance required in a disaster situation. If a high-priority process cannot tolerate degraded performance, a backup server equivalent to the primary server should be maintained. In other cases where a lower level of performance is acceptable, for example with an email system or a management reporting system, the applications could be run on virtual servers hosted on a small number of physical servers that are shared with other applications.
- Sufficient storage to restore essential data to continue operations and to accommodate new data generated while in disaster recovery mode—Once again, it is essential to prioritize. Not all business data is equally valuable. For example archived data from department-level data warehouses or operational data stores is probably not necessary until normal operations are fully restored. Rather than incur the cost of maintaining disaster recovery storage for all possible business data, maintain only as much as is needed for necessary operations.
- Office and data center space—This includes physical space for employees and equipment as well as critical infrastructure, especially power.

Proper planning will help identify essential business operations, levels of service required, and the minimal amount of physical space required in a disaster situation. This planning in turn will help control the cost of maintaining the physical requirements for disaster recovery.

Application Requirements

With the physical requirements attended to, the next step is planning on ensuring data and applications are up to date and ready to carry on business operations in the event of disaster. A basic question that must be answered is, How long can operations be down before there are significant, adverse consequences? Can your business, for example, continue to operate if data on backup tapes has to be restored to standby servers before the disaster recovery site becomes operational? Can the business recreate data that was created after the last offsite backup and before the disaster struck? If the answer is no to either question, replication and high-availability services should be considered.

Replication systems are used to copy data from primary servers to standby servers on an ongoing basis. Replication systems can efficiently capture changes to disks, copy the data to standby servers, and ensure data is duplicated offsite in a reasonably short period of time. This approach does require continuously operating hardware at both the primary and disaster recovery sites as well as sufficient bandwidth to keep the standby server up to date.

Replication services maintain the data but do not control the process of switching operations from the primary to the standby servers. High-availability systems perform this function. If rapid failover is needed, high-availability systems can be configured to monitor primary servers and automatically switch to the standby servers as soon as a failure is detected. As a general rule, the more speed and automation required in the failover process, the greater the cost. In situations where disaster recovery budgets are constrained and slower failovers are tolerable, a manual failover process can be used.

Disaster Recovery Service Providers

An alternative to maintaining a dedicated disaster recovery site is to use a third-party disaster recovery service. The business model of such operations is based on the idea of pooled risk. It is not likely that all customers would need to use standby servers at the same time, so they can more efficiently provide disaster recovery services than a typical business.

Maintaining continuous business operations in the event of disaster is one of the most difficult challenges in recovery management. Planning is a key to success. By prioritizing business operations and identifying operations that can tolerate degraded performance, business can find a balance between maintaining services and controlling the cost of disaster recovery services.

Summary

Recovery management entails a number of technical challenges. Scheduling and monitoring backups, choosing the right storage media, controlling the cost of offsite storage, keeping up with growing data volumes, and planning for disaster recovery all present technical issues. Understanding business drivers behind recovery management is necessary to properly prioritize business operations which in turn is essential to making the appropriate choices when confronting each of these challenges.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.