

# Determining Windows 2000 Storage Management Strategies



Network administrators who handle day-to-day network operations and who know the requirements of data management and storage systems need to become familiar with new Microsoft® Windows® 2000 Server features related to storage management. When you are planning your Windows 2000 deployment, it is recommended that you incorporate these new features into your storage management strategy. Disk Management, Removable Storage, Remote Storage, Windows Clustering, Distributed file system, Microsoft® Indexing Service, and other features can help you improve your storage management functions.

Considerations for selecting a data storage system, fault tolerance and backup strategies, and ways to improve your disaster recovery capabilities are also discussed in this chapter.

## **In This Chapter**

- Improving Your Storage Management Functions 699
- Managing Disk Resources 704
- Optimizing Data Management 712
- Enhancing Data Protection 722
- Improving Your Disaster Recovery Capabilities 724
- Planning Task List for Storage Management 728

### **Chapter Goals**

This chapter will help you develop the following planning documents:

- Storage configuration policy
- Disaster recovery plan
- Storage management plan

### **Related Information in the Resource Kit**

- For more information about using Removable Storage and Remote Storage, see “Data Storage and Management” in the *Microsoft® Windows® 2000 Server Resource Kit Server Operations Guide*.
- For more information about backup planning, strategies, and procedures, see “Backup” and “Planning a Reliable Configuration” in the *Server Operations Guide*.
- For more information about disaster recovery, see “Repair, Recovery, and Restore” in the *Server Operations Guide*.

## Improving Your Storage Management Functions

Improving your storage systems and their management is not only an important consideration when deploying Microsoft Windows 2000 Server but is also a critical part of any enterprise network infrastructure. With the enormous amount of data that must be protected in an enterprise environment, you need to be aware of the latest technologies so you can select the hardware and software that best meet your network needs.

Microsoft Windows 2000 provides several features for managing disk resources to enhance performance and protect data. These features include the following:

**Disk Management** for setting up and organizing disk storage systems.

**Removable Storage** for managing a new class of storage devices.

**Remote Storage** for moving unused files to remote storage.

The following Windows 2000 features can help you manage data more effectively:

**Windows Clustering** for easier management and higher availability of data and applications.

**File system improvements** for improving the performance, availability, security, and manageability of shared information and resources, including the NTFS file system and quota management.

**Distributed file system (Dfs)** for linking shares into a single namespace, making it easier to find and manage data.

**Indexing Service** for fast file searches by content and property.

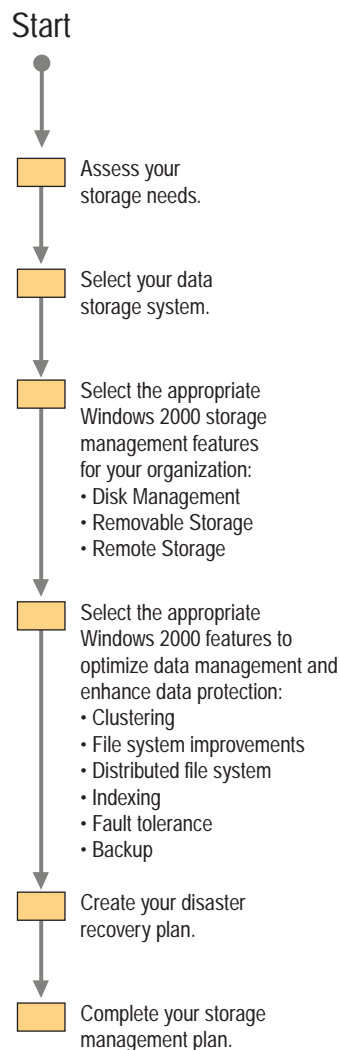
In addition to these features, Windows 2000 offers fault tolerance and Backup features to help you improve your data protection.

The sections that follow discuss these features in more detail. In addition to becoming familiar with Windows 2000 storage management features, you need to include a storage management plan in your deployment planning documents.

## Creating Your Storage Management Plan

You will have many issues to address when you create your storage management plan. Enterprise-size organizations should consider creating a storage management team to identify data storage needs and to create associated plans. Some organizations might find that the issues in this chapter are best addressed by a variety of teams, such as a backup and recovery team; a data management team; and a team handling storage issues. Each team begins by assessing the storage requirements of your organization and by creating a storage management strategy.

To help you create a storage management plan, consider using the steps shown in Figure 19.1.



**Figure 19.1** Process for Developing a Storage Management Strategy

## Assessing Your Storage Needs

Network storage solutions are rapidly becoming available as the number and size of enterprise networks increases. Every organization has different priorities for selecting media and methods for data storage. Some are constrained by costs, and others place performance before all other considerations.

As you assess your storage needs, you need to compare the possible loss of data, productivity, and business to the cost of a storage system that provides high performance and reliability. Consider the following needs before you develop your storage management strategy:

- Technologies that are the most cost-effective for your organization.
- Adequate storage capacity that can easily grow with your network.
- The need for rapid, 24-hour access to critical data.
- A secure environment for data storage.

When looking for the most cost-effective solution, you need to balance the costs of purchasing and maintaining hardware and software with the consequences of a disastrous loss of data. Costs include:

- Initial investment in hardware, such as tape and disk drives, power supplies, and controllers.
- Associated media such as magnetic tapes and compact discs.
- Software, such as storage management tools and a backup tool.
- Ongoing hardware and software maintenance costs.
- Staffing.
- Training in how to use new technologies.
- Off-site storage facilities.

Compare these costs against:

- Replacement costs for file servers, mail servers, or print servers.
- Replacement costs for servers running applications, such as Microsoft® SQL Server™ or Microsoft® Systems Management Server (SMS).
- Replacement costs for gateway servers running Routing and Remote Access Service, Microsoft® SNA Server, Microsoft® Proxy Server, or Novell NetWare.
- Workstation replacement costs for personnel in various departments.
- Replacement costs for individual computer components, such as a hard disk or a network card.

Another important factor to consider when you select a storage system is speed of data recovery. If you lose the data on a server, how fast can you reinstate that data? How long can you afford to have a server (or an entire network) down before it begins to have a serious impact on your business?

Storage technology changes rapidly, so it is best to research the relative merits of each type before you make a purchasing decision. The storage system you use needs to have more than enough capacity to back up your most critical data. It should also provide error detection and correction during backup and restore operations.

## Selecting a Data Storage System

Answer the following questions to determine a storage system that best meets your needs:

### **How much data do you currently need to store?**

If a very large amount of data must be stored, a tape-based storage system might be your best choice. The media cost per megabyte is significantly lower for tape than it is for many other types of storage media.

### **What are your projected data-storage needs?**

The storage needs of many organizations double every year. Consider purchasing a larger storage system than is necessary to satisfy present requirements or a scalable system that you can expand as your needs grow. To evaluate how this might affect your situation, compare your data storage from several years ago to your current levels, and use this growth to estimate your future needs.

### **How many users or applications simultaneously access the data storage system?**

Most vendors offer multidrive systems that allow several drives to be accessed at once. In this way, multiple users or applications can simultaneously access the system without affecting performance.

### **How important is data-access time?**

If your library is used primarily to access data for real-time use, then this is your most important issue. CD-ROM disc-based solutions usually work best if data-access time is your primary concern because the random-access capability of CD-ROM discs reduces seek time to a minimum. Data-access time has two parts: seek time and transfer rate. The disadvantages of this solution are in speed and cost: the data transfer rate can be slower than tape-based systems unless the latest high-speed drives are used; and the cost per megabyte is greater than tape media.

**How important is the data-transfer rate?**

If your data storage system is used primarily for archiving and backing up data, then the data-transfer rate is your most important issue. If this is the case, then a tape-based solution might be best because the data-transfer rate of tape drives is approximately ten times the speed of CD-ROM disc-based drives. Tape-based systems are also less expensive per megabyte. The disadvantage of tapes is that file-access time is increased by as much as several minutes because file access is linear.

**What is your budget?**

Again, before you decide how much you can afford to spend, factor in the potential costs of lost or corrupted data and downtime due to problems with unreliable hardware. If the type of data that you store is important to your organization, these risks might not be worth the savings gained by purchasing an inexpensive solution. Also, consider the whole cost. Certain hardware can be relatively inexpensive to acquire but the cost per megabyte is expensive. This is generally true for CD-ROM disc-based systems. It is also common to spend more money over time on the media than on the initial hardware.

To help you select a storage system, try creating two or more models that present different hardware and software solutions for varying degrees of storage capacity and data protection. Be sure to project for planned growth.

Table 19.1 shows the *relative* capabilities of possible hardware and software storage solutions: a 5 indicates this is the best available solution; a 1 indicates this is your least desired solution. Use the table to help you choose types of storage for your organization.

**Table 19.1 Selected Hardware and Software Storage Solutions**

<b>Solution</b>	<b>Availability</b>	<b>Response Time</b>	<b>Capacity</b>	<b>Support for Multiple Users</b>
CD-ROM/DVD-ROM	4 <sup>1</sup>	3	2-3	2
CD-ROM library	4	2	5	5
CD-ROM drive array	5	4	4	4
Dfs	5	3-4	5	5
Disk	3	4	3	3
Disk mirroring with two controllers (Duplex)	5	4	2-3	2

*(continued)*

**Table 19.1 Selected Hardware and Software Storage Solutions** *(continued)*

<b>Solution</b>	<b>Availability</b>	<b>Response Time</b>	<b>Capacity</b>	<b>Support for Multiple Users</b>
Disk stripe set	1	5	4	4
Disk stripe set with parity	4	3	3-4	4
Tape	3	2	4	1
Tape library	3	1	5	4

1 5 – High/1 – Low

After you have assembled a storage management team, identified your storage requirements, and determined your budget, you need to evaluate the storage capabilities of Windows 2000.

## Managing Disk Resources

Windows 2000 offers several storage features to help you store and manage data, including Disk Management, Removable Storage, and Remote Storage. The subsections that follow provide an introduction to these features.

## Disk Management

The Disk Management snap-in to the Microsoft Management Console (MMC) is a tool for managing disk storage systems. Wizards guide you through creating partitions or volumes and initializing or upgrading disks. New key features of Windows 2000 Server Disk Management include the following:

**Online Disk Management** You can perform most administrative tasks without shutting down the system or interrupting users. For example, you can create various partition layouts and choose protection strategies, such as mirroring and striping, without restarting the system. You can also add disks without restarting. Most configuration changes take effect immediately.

**Remote Disk Management** As an administrator, you can manage any remote (or local) computer that runs Windows 2000.

Figure 19.2 shows some of the View menu options you can select in Disk Management.

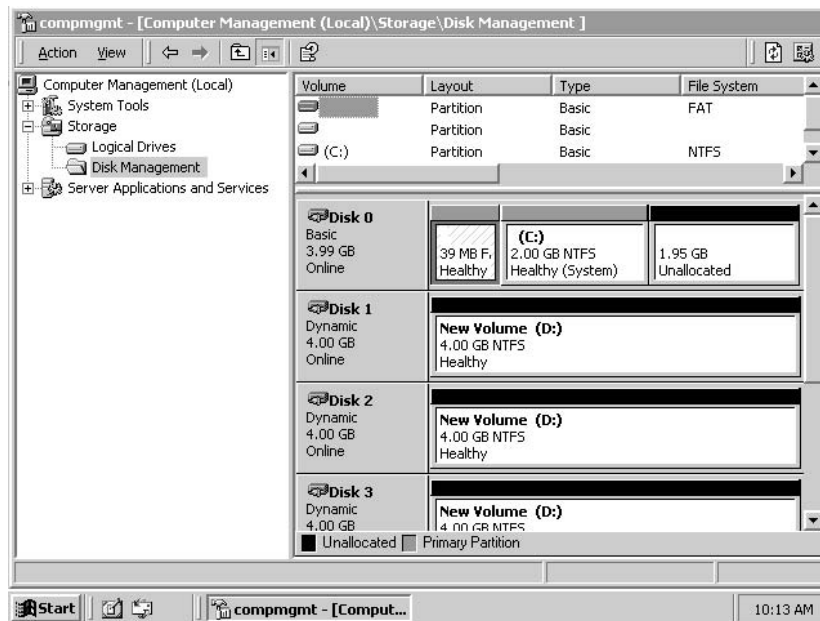


Figure 19.2 Disk Management MMC Snap-In

## Basic and Dynamic Storage

There are two types of disk storage available with Windows 2000: basic or dynamic. *Basic storage* supports partition-oriented disks. A basic disk can hold primary partitions, extended partitions, and logical drives. Basic disks might also contain spanned volumes (volume sets), mirrored volumes (mirror sets), striped volumes (stripe sets), and redundant array of independent disks or RAID-5 volumes. In Microsoft® Windows NT® version 4.0 or earlier, RAID-5 was known as a stripe set with parity. If you want computers to access these volumes, and if those computers run Windows NT 4.0 or earlier, Microsoft® Windows® 98 or earlier, or Microsoft® MS-DOS®, you need to create basic volumes.

*Dynamic storage* supports new volume-oriented disks and is new with Windows 2000. It overcomes the restrictions of partition-oriented disk organization and facilitates multidisk, fault-tolerant disk systems. With dynamic storage, you can perform disk and volume management without restarting the operating system. On a dynamic disk, storage is divided into volumes instead of partitions. A volume consists of a portion or portions of one or more physical disks in any of the following layouts: simple, spanned, mirrored, striped, and RAID-5 volumes. Dynamic disks cannot contain partitions or logical drives, and cannot be accessed by MS-DOS or Microsoft® Windows® 98 and earlier versions. You can use dynamic storage to set up a fault-tolerant system by using multiple disks.

When you attach a new disk to your computer, you need to initialize the disk before you can create volumes or partitions. When you initialize the disk, select dynamic storage if you want to create simple volumes on the disk or if you plan to share the disk with other disks to create a spanned, striped, mirrored, or RAID-5 volume. Select basic storage if you want to create partitions and logical drives on the disk.

Table 19.2 shows tasks that you can perform on basic and dynamic disks by using Disk Management.

**Table 19.2 Tasks for Basic and Dynamic Disks**

Tasks	Basic Disk	Dynamic Disk
Create and delete primary and extended partitions.	X	
Create and delete logical drives within an extended partition.	X	
Format and label a partition and mark it as active.	X	
Delete a volume set.	X	
Break a mirror from a mirror set.	X	
Repair a mirror set.	X	
Repair a stripe set with parity.	X	
Upgrade a basic disk to a dynamic disk.	X	
Create and delete simple, spanned, striped, mirrored, and RAID-5 volumes.		X
Extend a volume across one or more disks.		X
Add a mirror to or remove a mirror from a mirrored volume.		X
Repair a mirrored volume.		X
Repair a RAID-5 volume.		X
Check information about disks, such as capacity, available free space, and current status.	X	X
View volume and partition properties such as size.	X	X
Make and change drive-letter assignments for hard disk volumes or partitions and CD-ROM devices.	X	X
Create <i>volume mount points</i> .	X	X
Set or verify disk sharing and access arrangements for a volume or partition.	X	X

## Volume Management

Windows 2000 includes significant improvements in the architecture of volume management. Volume management includes the processes that create, delete, alter, and maintain storage volumes in a system. The new architecture improves the manageability and recoverability of volumes in an enterprise environment.

A Logical Disk Manager (LDM) has been introduced to the architecture to extend fault tolerance functionality, to improve system recovery, to encapsulate volume information so that disks can be easily moved, and to provide improved management functionality. This service is responsible for volume creation and deletion, fault tolerance features (RAID), and volume tracking. You use the Disk Management snap-in to manage local and remote volumes.

Volume management has the following features:

- You can create any number of volumes in the free space on a physical hard disk or create volumes that span two or more disks.
- Each volume on a disk can have a different file system, such as the file allocation table (FAT) file system or the NTFS file system.
- Most changes that you make to your disk are immediately available. You do not need to quit Disk Management to save them or restart your computer to implement them.

## Volume Mount Points

As part of Disk Management, you can create volume mount points. Volume mount points provide you with a quick way to bring data online and offline. They are file system objects in the Windows 2000 internal namespace that represent storage volumes. When you place a volume mount point in an empty NTFS directory, you can graft new volumes into the namespace without requiring additional drive letters. An example of how you might use volume mount points is to have a computer with a single drive and volume formatted as C and to mount a disk as C:\Games.

Some possible uses for volume mount points include:

**To provide additional space for programs** For example, you mount a disk as C:\Program Files. Then, when you need additional disk space, you add a disk to the system and span it with the disk at C:\Program Files.

**To create different classes of storage** For example, create a stripe set for performance and mount it as C:\Scratch; and create a mirror set for robustness and mount it as C:\Projects. Users will see the directories normally, but their scratch directory will be fast, and their projects directory will be protected by a mirror set.

**To create multiple mount points for a volume** For example, a volume is mounted as C:\Games and C:\Projects. Be aware that nothing prevents cycles in the namespace. If you mount a volume as D and also as D:\Docs, because D is mounted underneath itself, it creates a cycle in the namespace. Applications that do enumeration get into an endless loop on this volume.

Volume mount points are robust against system changes that occur when hardware devices are added to or removed from a computer. You are no longer limited to the number of volumes you can create based on the number of drive letters.

## Disk Defragmentation

Another Disk Management feature is the Disk Defragmenter. You can use this tool to locate files and folders that have become fragmented and to reorganize clusters on a local disk volume. Disk Defragmenter organizes clusters so that files, directories, and free space are physically more contiguous. As a result, your system can gain access to your files and folders and can save new ones more efficiently. If you have considerable fragmentation, Disk Defragmenter can improve your overall system performance significantly in relation to disk input/output (I/O).

The Disk Defragmentation feature decides where files should be located on the disk, but NTFS and FAT move the clusters around.

You can use this tool with disk volumes that are formatted for FAT16, FAT32, or NTFS.

## Considerations for Using Dynamic Storage

Consider the following when creating volumes:

- Dynamic storage uses a volume-oriented scheme for disk organization. Windows NT Server is not compatible with dynamic disks.
- You can use the Windows 2000 Setup program to configure disk space while upgrading to Windows 2000 Server.

---

**Note** You can create new volumes and partitions on unallocated portions of the disk without losing data on existing volumes. However, if you plan to change your volume topology, you have to backup your data because making changes to existing volumes erases all existing data.

---

- You can configure the internal hard disk on a new computer during initial setup when you load the Windows 2000 Server operating system software. You can use Disk Management to make changes to the disk after installation.

For more information about managing disks, see “Disk Concepts and Troubleshooting” and “Data Storage and Management” in the *Microsoft® Windows® 2000 Server Resource Kit Server Operations Guide*.

## Removable Storage

The Removable Storage system is a new technology that enables multiple applications to share local libraries and tape or disk drives, thus providing better storage management functionality. With Removable Storage, you can use stand-alone storage devices; manage online media libraries and robotic changers; and track your removable tapes and disks. Stand-alone devices include CD-ROM, DVD-ROM, tape (4mm, DLT, 8mm, and others), and high-capacity disk drives.

Removable Storage also controls removable media within a single server system. In addition, it performs functions in conjunction with Backup and Remote Storage. A key aspect of Removable Storage is the ability of applications to create media pools that are owned and used by the application.

Storage devices are most often connected to systems by using small computer system interface (SCSI) adapters or integrated device electronics (IDE) interfaces, such as those used with most hard drives. New technologies that provide simpler use and faster throughput, such as Fiber Channel, IEEE 1394, and Intelligent I/O (I2O), are also being used with increasing frequency. Stand-alone devices are most commonly used with single user systems.

Libraries contain multiple drives using CD-ROM, DVD-ROM, magneto-optic (MO) discs, or tape. These are accompanied by robotic controls that provide extensive automation for managing individual pieces of media or storage. The capacities range from a small, three-disc CD-ROM autochanger to a tape or a disc library used with sophisticated applications inside a large corporation. Libraries are most commonly used with servers but are increasingly being attached to single user systems.

Tasks you can perform using Removable Storage include the following:

- Tracking online and offline media
- Mounting and dismounting media
- Inserting and ejecting media to and from a library
- Viewing the state of media and libraries
- Creating media pools and setting media-pool properties

- Setting security parameters for media and media pools
- Performing library inventories

---

**Note** Your backup software must be compatible with Removable Storage to make use of these features.

---

## Remote Storage

Remote Storage is the hierarchical storage management system for Windows 2000 Server. With Remote Storage, you use the Remote Storage Manager MMC snap-in to move unused files to a tape library. By regularly migrating files, you can increase the amount of free space on a disk. From the user's perspective, the migrated files remain active, but they take longer to access.

The storage hierarchy has two levels. The highest level is called local storage. It is made up of the local NTFS volumes of a computer running Remote Storage on a Windows 2000 server. The local disk volumes that are under Remote Storage control are called *managed volumes*.

The lower level of the storage hierarchy, called remote storage, stores the data that has been copied from local storage to an online library or additional storage device.

When the amount of free space on a local volume drops below the level you need, Remote Storage truncates data from local files that were previously copied to remote storage, and by doing so, provides free disk space. When data is moved, it leaves a file marker allowing you to access that file. Remote Storage manages movement of data according to guidelines set by the administrator for each local storage volume. You can set a schedule for moving files from specific volumes and set criteria and rules for the files that are to be moved. More specifically, you can:

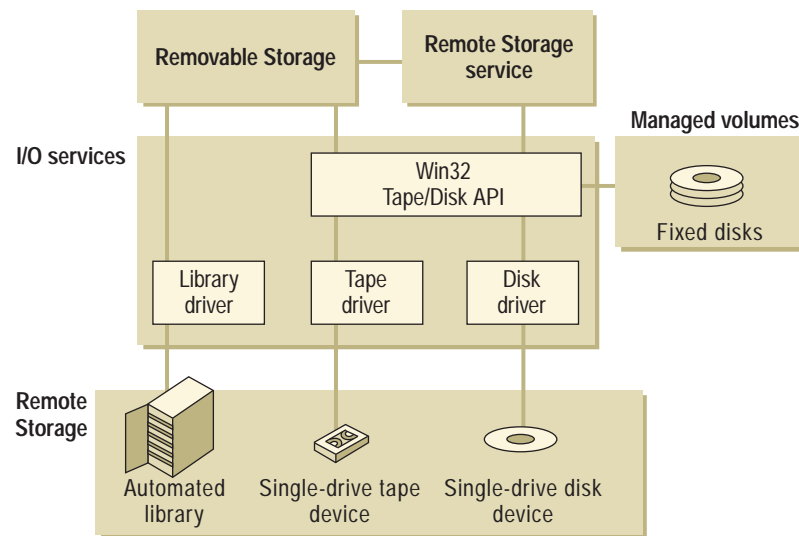
- Allocate and configure remote storage devices and media.
- Set Remote Storage system-wide feature options.
- Configure volume management settings for Remote Storage managed volumes.
- View information about Remote Storage activity.
- Recover from media disasters.
- Create and submit jobs.

Because removable tapes in a library are less expensive per megabyte than hard disks, this is an economical way to provide both maximum data storage and optimal local disk performance.

**Note** Your backup and virus scanning software must be compatible with Remote Storage. Administrators need to ensure that volume-wide file operations are done before activating Remote Storage so that you do not have to move everything back to the disk. Backup reads the data directly from the tape.

## Relationship of Remote Storage and Removable Storage

Remote Storage uses Removable Storage for copying data to online libraries that contain removable media. Figure 19.3 provides an overview of the relationship of these storage systems and various storage devices.



**Figure 19.3** Relationship of Remote Storage, Removable Storage, and Storage Devices

## Considerations for Using Remote Storage

Using Remote Storage provides the following advantages:

- Virtual expansion of local storage space by using lower cost remote storage.
- Transparent automatic access to data in remote storage.
- Automation of the labor-intensive overhead associated with daily manual data management operations.
- Centralized sharing of remote storage for multiple volumes.

Remote Storage is not a replacement for backup because only one instance of data exists. It is important to regularly backup the volume. Backup is integrated with Remote Storage so you do not have to move everything back to the disk; backup reads the data directly from the tape.

## Optimizing Data Management

Windows 2000 features that can help you manage data more effectively in an enterprise environment include:

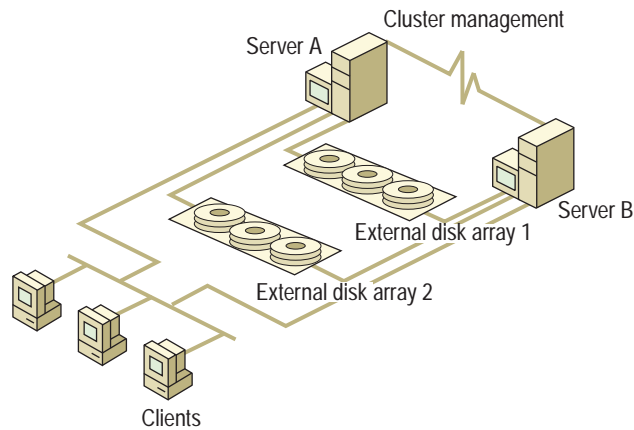
- Windows Clustering
- NTFS file system
- Quota management
- Distributed file system (Dfs)
- Indexing Service

## Windows Clustering

Consider clustering in your enterprise network storage strategy when you need greater availability with simplified management. Clustering reduces downtime by providing an architecture that keeps systems running in the event of a single server failure.

With Windows Clustering, you can connect two or more servers to form a cluster of servers that work together as a single system. Each server is called a node; each node can operate independently of the other nodes in the cluster. The Windows 2000 built-in clustering capability is based on open specifications, industry-standard hardware, and ease-of-use requirements.

Each node has its own memory, system disk, operating system, and subset of the cluster's resources. With a process called failover, if one node fails, the other node takes ownership of the failed node's resources. The cluster server then registers the network address for the resource on the new node so that client traffic is routed to the available system, which now owns the resource. When the failed resource is later brought back online, you can configure the cluster server to redistribute resources and client requests appropriately. A standard Windows 2000 cluster setup is shown in Figure 19.4.



**Figure 19.4** Typical Two-Node Cluster Setup

The Clustering Service provides the following advantages:

**Common Administration** With Cluster Administrator MMC snap-in, you administer a cluster as a single system. Also, a client computer interacts with a cluster as though it is a single server.

**Load Balancing** Within the cluster, you can manually balance processing loads or unload servers for planned maintenance without taking data and applications offline.

**High Availability** Clustering provides high availability by automatically recovering critical data and applications from many common types of failure. If a node in the cluster fails, Windows Clustering detects the failure and performs recovery of the processes active at the time of the failure. The failure of one node in the cluster does not affect the other node.

If clustering is not already part of your network and if you want to deploy Windows Clustering, you need to address several issues in your Windows 2000 Server deployment planning phase with regard to setting up a cluster environment. For more information about planning your cluster environment, see “Windows Clustering” in the *Microsoft® Windows® 2000 Server Resource Kit Distributed Systems Guide*, and also see Windows 2000 Advanced Server Help.

---

**Note** For clustering solutions, use only certified configurations that are documented on the hardware compatibility list (HCL), which you can access online. For more information about this list, see the Microsoft Windows Hardware Compatibility List link on the Web Resources page at <http://windows.microsoft.com/windows2000/reskit/webresources>.

---

## Considerations for Using Clustering in Your Storage Strategy

Consider the following advantages of incorporating a cluster environment into your storage planning:

- Clustering provides high availability without data replication, thereby providing data consistency without greatly impacting storage requirements and network traffic volumes.
- Clustering provides the ability to easily recover from software failures.
- Servers share a multiported disk array. Hardware RAID controllers provide the best performance to the external disk array.
- Clustering provides high availability of data but does not protect data integrity.

## File System Improvements

Windows 2000 supports the NTFS file system and two file allocation table (FAT) file systems: FAT16 and FAT32.

FAT is for small disks and simple folder structures. FAT16 is included with Windows 2000 because it maintains an upgrade path for earlier versions of Windows-compatible products and it is compatible with most non-Microsoft operating systems. FAT32 supports volumes larger than those handled by FAT16, and was first available in Microsoft® Windows® 95. Windows 2000 supports FAT32 file systems.

### NTFS

The version of NTFS used in Windows 2000 provides performance, reliability, and features that are not available with FAT file systems. The NTFS data structures allow you to take advantage of new features in Windows 2000, such as Microsoft® Active Directory™ directory service, change and configuration management, reparse points (directory junctions and volume mount points), sparse file support, object IDs, extended property attributes, change journal, and many of the new storage enhancements.

In Windows 2000, NTFS data structures have been updated to enable many new features. For existing NTFS volumes, the upgrade to the new version of NTFS occurs when you install Windows 2000. You can convert FAT16 and FAT32 volumes to this format at any time.

Formatting your Windows NT partitions with NTFS instead of FAT allows you to use recoverability and compression features that are available only with NTFS. Also, formatting your volumes with NTFS instead of FAT provides faster access speed and additional capabilities for file and folder security.

---

**Important** The version of NTFS used in Windows 2000 cannot be natively recognized by earlier versions. In dual-boot systems, where a Windows NT 4.0 installation needs to read an NTFS volume that was created or upgraded by Windows 2000, the Windows NT 4.0 installation will require support (Service Pack 4 or later).

---

## Quota Management

Disk quotas are a new feature with the version of NTFS used in Windows 2000. Disk quotas provide more precise control of network-based storage. You can use disk quotas to monitor and limit disk space use on a per user, per volume basis.

The first time users try to store data on a volume, they are automatically entered in the quota table and assigned a default quota value. This means that the administrator does not have to enter a quota setting for each user.

Users are charged for the files they own. For example, each user's folder on \\Marketing\Public is limited to 5 megabytes (MB) of disk space. If users copy 5 MB of files to their folder, they cannot then copy or create any more files on this or any other folder on \\Marketing\Public. They can, however, move or delete the files. Disk space is not charged to users if they modify an existing file owned by someone else. Keep in mind, however, that some applications, for example Microsoft® Office, change the owner of a document to the user who last edited the document. Quota settings are independent across volumes; that is, the quota on drive C does not affect the quota on drive D.

You can use the quota feature of the Disk Management MMC snap-in to:

- Enable or disable quotas on a disk volume.
- Prevent users from using more disk space when their quota limit is exceeded.
- View quota information for each user of the volume.
- Set the default quota warning threshold and quota limit assigned to new volume users.
- Block additional disk allocations and log an event when a user exceeds a specified disk space limit. Users can read, delete, and edit files as long as they do not attempt to allocate more disk space.

You can set both threshold and hard quota limits. When you enable quotas, you can set two values:

**Quota limit** Specifies the maximum amount of disk space a user is allowed to use.

**Quota warning threshold** Specifies a value at which the administrator is alerted that a quota limit is being approached. This is in the form of an event message.

As administrator, you can specify that events be automatically logged when users exceed warning thresholds and quota limits. For example, you can set a user's disk quota limit to 50 MB, and the quota warning level to 45 MB. If the user stores more than 45 MB of files on the volume, the quota system logs a system event.

You have the option of denying disk space to users who attempt to exceed their quota limit. If you select this option, users cannot write additional data to the volume without first deleting or moving some existing files from the volume. NTFS displays an "out of disk space" error message if the user tries to allocate beyond their quota limit.

Windows 2000 includes disk quota support for the following:

- Policies for wide-scale remote management of disk quotas.
- Improved support for finding all files owned by a particular user.

---

**Note** Windows 2000 Server supports disk quotas only for the volumes formatted as NTFS.

---

When creating your storage management strategy, consider the following advantages of using disk quotas:

- Monitoring disk space use on a per user and per volume basis allows for better disk resource planning.
- Limiting disk space allows you to manage your storage resources more effectively by encouraging users to regularly delete unnecessary files.
- Using disk quotas can be an effective way to decrease backup media costs and restoration times.

## Distributed File System

Microsoft Distributed file system (Dfs) is Windows 2000 Server software that makes it easier for you to find and manage data on your enterprise network. Dfs provides mapping and a uniform naming convention for collections of servers, shares, and files. Dfs adds the capability of organizing file servers and their shares into a logical hierarchy, making it considerably easier to manage and use information resources.

With Dfs, you can create a single directory tree that includes multiple file servers and file shares in a group, division, or enterprise. Any Windows 2000 server can host a Dfs root or Dfs volumes. A Dfs root is a local share that serves as the starting point and host to other shares. A network can host many individual Dfs volumes, each having a distinct name. A Dfs topology is a single Domain Name System (DNS) namespace. You can use a single topology or multiple Dfs topologies to distribute your organization's shared resources.

Dfs functionality is integrated with Active Directory; the Dfs topology is published to Active Directory. Because changes to a domain-based Dfs topology are automatically synchronized with Active Directory, you can always restore a Dfs topology if the Dfs root is offline for any reason. Computer-based Dfs stores the topology in the registry. Dfs has the following features:

- Provides a simplified view of network shares that can be customized by the administrator.
- Allows Microsoft® Windows® 95 and Windows 98 clients to access shares by using server message block (SMB) protocol.
- Supports mounting replicas of network shares for load balancing and better data availability.

Active Directory further optimizes network use by redirecting Active Directory-enabled clients to a Dfs share point within the client site.

- Integrates with the File Replication service (FRS) to permit optional replication of read/write data between multiple shares.
- Allows users to log on just once for multiple access.

You can access a Dfs volume by using a uniform naming convention (UNC) name. Although UNC names can be used, in most cases users will find it easier if they substitute a drive letter. For example, note the physical locations in relation to the logical paths shown in Table 19.3.

**Table 19.3 Accessing a Dfs Volume**

Dfs Logical Path	Physical Location	Description	Mapped Drive Path
\\MS Server\Root	\\MS Server\Root	Root share	X
\\MS Server\Root\Users	\\MS Users1\Employees	Junction to employee directories	X:\Users

*(continued)*

**Table 19.3** Accessing a Dfs Volume (*continued*)

Dfs Logical Path	Physical Location	Description	Mapped Drive Path
\\MS Server\Root\Private\ JaneD	\\Legal\Data\JaneD	Junction to JaneD's computer	X:\Private\JaneD
\\MS Server\Root\Private\ SusanY	\\Human Res\SusanY	Junction to SusanY's computer	X:\Private\SusanY

Since Dfs maps the physical storage into a logical representation, the physical location of data becomes transparent to users and applications. Dfs eliminates the need for users to know where information is physically stored. Because users do not need to know the name of a server or share, you can physically move user information to another server without reeducating users about how to find their data, thereby improving file management.

## Considerations for Using Dfs in Your Storage Strategy

Consider the following advantages of using Dfs shares as part of your storage planning:

- Active Directory replicates the Dfs topologies for all domain-based Dfs topologies to each Dfs root server. This distributes the load on participating servers and implements fault tolerance for the Dfs root.
- Multiple servers can host domain-based Dfs roots and alternates. If a root fails, Dfs detects the failure and another server acquires the root. This failover process increases data availability.
- Multiple copies of shares on separate servers can be mounted under the same logical Dfs name. This provides alternate locations for accessing data, thereby providing for load balancing and better data availability.
- Multiple copies of shares also allow administrators to perform preventative maintenance on servers. A server hosting one replica can be taken offline without impacting users since Dfs automatically routes requests to a replica that is online.
- Dfs ensures that users go to the closest replica by distributing copies of your files by site. This reduces the load on the wide area network (WAN).
- Location transparency eases the burden of upgrading to new servers by allowing additional storage to be published in subdirectories.

Consider implementing Dfs if any of the following conditions exist in your organization:

- Users that access shared resources are distributed across a site or sites.
- Most users require access to multiple shared resources.
- Users require uninterrupted access to shared resources.
- Load balancing for your network could be improved by redistributing shared resources.
- Your organization has data that is stored on multiple network shares.

For more information about designing a Dfs tree, see Windows 2000 Server Help.

## Indexing Service

Microsoft Indexing Service makes it easier for users to search for data on client computers and servers. Indexing Service scans files on Windows 2000 servers and client computers and builds content and property indexes that dramatically improve search capability and performance. When the service is running, users can search for words and phrases in thousands of files in just a few seconds.

Indexing Service has the following features:

- Searches by content (for example, searches all files containing “revenue projections”).
- Searches by document properties (for example, searches all files where AUTHOR contains “Sarah”).
- Searches with Boolean operators (for example, AND, OR, NOT).
- Uses a free-text search, which allows users to enter any combination of words without having to learn a particular search syntax.
- Can index volumes on the local computer and also network shares, including NetWare and UNIX servers.
- Provides secure query results.

Returns only the documents that users are allowed to read. Uses standard Windows 2000 access control lists (ACLs).

- Integrates with NTFS for better performance and reliability.
- Integrates with Internet Information Services (IIS) to provide a search capability for Internet and intranet Web sites.
- Can create customized search forms and user interfaces by using OLE-DB or Microsoft® ActiveX® Data Objects (ADO) scripting.

- Indexes a variety of file formats, including Microsoft® Office 97, Microsoft® Office 2000, text files, and HTML pages.
- Integrates with Windows 2000 user interface and Windows Explorer.
- Provides easy administration by using a Microsoft Management Console snap-in.

When Indexing Service is running on a system, it monitors the server for file modification. When files are modified, they are opened and their contents indexed. Opening files is done by a low priority background process so that general server performance is minimally impacted. In addition, when running on NTFS, Indexing Service uses a number of NTFS advanced features to minimize overall system overhead.

---

**Note** When you first run the service, it must build its indexes from scratch. This involves scanning all the files on the volume. Initial index construction accesses the disk heavily until the indexes have been built. After the indexes have been generated, only incremental updates are needed as files are modified, so further updates are virtually unnoticeable. In all cases, index update is a low priority task and will pause if server resources are needed for other operations.

---

To search for documents, users just need to select Search for Files and Folders in Windows Explorer or from the Start menu. This brings up a search form allowing users to enter the words for their search. If Indexing Service is running on a file server, users can search network shares efficiently because the search is executed on the server and only the search results are returned across the network.

## Integration with Windows 2000 Components

Indexing Service is integrated with many other Windows 2000 components for performance and reliability. The service also uses NTFS features such as bulk ACL processing for much faster security checks before returning search results. It also uses NTFS sparse files to optimize indexes without consuming extra disk space. The service uses the NTFS change log to monitor the volume for file modifications. In this way, the service does not scan the entire volume repeatedly for changes like many other search engines. Instead, when a file is modified, only that particular file is scanned and indexed.

Indexing Service also understands that files can be migrated by Remote Storage. It does not forcibly recall files to index them. It does not rescan if a file has been migrated to secondary storage. This means that even if files have been migrated to tape, users can still search for them. This is ideal if you use Remote Storage to maintain an archival store for documents.

Indexing Service can be switched into a read-only mode. This allows the administrator to back up the indexes. When in read-only mode, the service continues to execute queries but does not update the indexes. The indexes are guaranteed to be consistent and stable so a valid backup can be made. After backing up, you can restore the service to normal operations, and any modifications made to files during the backup are processed normally.

Finally, the full-text indexing engine used by Windows 2000 is compatible with full-text indexing features of Microsoft® SQL Server™ version 7.0. By using the distributed query processor in SQL Server, you can specify queries using structured query language (SQL) and execute them simultaneously against the file system and a database.

## Considerations for Using Indexing Service in Your Storage Strategy

Consider the following advantages of using indexing as part of your storage planning:

- Users can easily and quickly find the documents they need on file servers and Web servers.
- Most users do not need to learn a search syntax, although a powerful search language is available for advanced users.
- A single file server and index can satisfy queries for multiple network shares, including files on non-Windows file servers.
- Performance is improved and system load is reduced due to tight integration with the Windows 2000 infrastructure.
- A secure search engine guarantees users cannot find documents that they are not allowed to view or read.
- User interface is easily customized using OLE and ADO programming interfaces.

Consider implementing Indexing Service if any of the following conditions exist in your organization:

- Users cannot find documents on servers or forget the location of documents.
- File servers contain many hundreds or thousands of documents making browsing for particular documents unwieldy or impossible.
- You need to provide search capabilities for a Web site.

## Enhancing Data Protection

In an enterprise network, you use a combination of strategies to protect your data. Using Backup and Windows 2000 fault-tolerance features are two ways you can improve your data protection.

## Fault Tolerance

Fault tolerance is the ability of a system to continue functioning when part of the system fails. Fault tolerance combats problems such as disk failures, power outages, or corrupted operating systems, which can impact startup files, the operating system itself, or system files. Windows 2000 Server includes fault-tolerant features.

Although the data is always available and current in a fault-tolerant system, you still need to make tape backups to protect the information about your disk subsystem against user errors and natural disasters. Disk fault tolerance is not an alternative to a backup strategy with off-site storage.

Fault-tolerant disk systems are standardized and categorized in six levels, known as RAID level 0 through level 5. Each level offers a specific mix of performance, reliability, and cost.

## Disk Management

Windows 2000 Disk Management includes RAID levels 1 and 5:

### **Level 1: Mirrored volumes (mirror sets in Windows NT 4.0)**

Mirrored volumes provide an identical copy for a selected volume. All data that is written to the primary volume is also written to a secondary volume or mirror. If one disk fails, the system uses data from the other disk. Because each file is stored in two locations, you need twice your usual storage space to implement this.

### **Level 5: RAID-5 volumes (striping with parity)**

RAID-5 volumes share data across all the disks in an array. The system generates a small amount of data, called parity information, that is used to reconstruct lost information in case a disk fails. RAID 5 is unique because it writes the parity information across all the disks. If a disk fails, data redundancy is achieved by arranging a data block and its parity information about different disks in the array. This level requires a minimum of three disks. As more disks are added to a RAID-5 set, the amount of overhead decreases from the maximum of 50 percent (that is, three disks are required to store the data normally on two disks). However, the benefits of having many disks in a RAID-5 set drops off when seven or more disks are used in the set.

## Selecting a RAID Strategy

RAID strategies include hardware and software solutions. Choosing between RAID-1 and RAID-5 volumes depends on your computing environment. Consider the following when selecting a RAID strategy:

- When compared to RAID-5 volumes, a mirrored volume implementation has a lower entry cost, requires less system memory, provides better overall performance, and does not show performance degradation during a failure. However, its cost per megabyte is higher than that for RAID-5 volumes.
- A software RAID-5 volume implementation has better read performance and a lower cost per megabyte, but it requires more system memory and loses its performance advantage when a disk in the array is missing.
- Hardware or software RAID-5 volumes are a good solution for data redundancy in a computing environment in which most activity consists of reading data. For example, you might want to use a RAID-5 volume on a server that is used to maintain all copies of the programs for your site. It enables you to protect the programs against the loss of a single disk in the striped volume. In addition, the read performance improves due to concurrent reads across the disks that make up the RAID-5 volume.
- In an environment in which frequent updates to the information occur, it might be better to use mirrored volumes. However, you can use a RAID-5 volume if you want redundancy and if the storage overhead cost of a mirror is prohibitive.

## Backup

The Backup program helps you protect data from accidental loss due to hardware or storage failure. Using Backup, you can create a duplicate copy of the data on your hard disk and archive the data on another storage device such as a hard disk or a tape. You can back up data to a wide variety of removable, high-density storage media. Also, Backup is integrated with Remote Storage for archiving.

With Backup wizards you can:

- Make an archived copy of selected files and folders on your hard disk.
- Schedule regular backups to keep your archived data up-to-date.
- Restore the archived files and folders to your hard disk or any other disk you can access.
- Back up Active Directory. You can put a copy of Active Directory store on media for off-site storage.
- Back up offline Remote Storage data, registry settings, and any other mount point data.

## A Data Protection Strategy for Enterprise Networks

Consider the following backup and fault-tolerance strategies when creating your data protection policies:

- Back up an entire volume in case a disk fails. It is more efficient to restore the entire volume in one operation.
- Always back up the directory services database on a domain controller to prevent the loss of user account and security information.
- For your critical computers, you can implement a software mirror of two separate hardware-controlled RAID arrays. With this configuration, if either a disk or an entire array fails, operations can continue.
- In the event that a computer running Windows 2000 Server fails, you should have a spare computer with Windows 2000 Server already installed to which you can move the data disks.

### Considerations for Designing a Fault-Tolerant Storage System

Some points to consider when planning your storage strategy include the following:

- In general, you only need to use fault-tolerant configurations for information that you must have readily available in case of hardware failure or unrecoverable disk errors should the primary data source go offline for any reason.
- If you have applications on a single computer running Windows 2000 Server, you only need to run them on a fault-tolerant volume if you cannot tolerate their unavailability for the amount of time it takes you to restore the applications from a backup.
- You need to back up the application volume any time you install a new application or change default settings for an application.
- If space is a consideration, you can format your application volume with the NTFS file system and use NTFS compression for folders and files on the volume.

### Improving Your Disaster Recovery Capabilities

Because a computer or site disaster can overcome even the best data protection strategies, you need to have a system disaster recovery plan. A disaster includes anything from not being able to start a computer to the destruction of a network when a natural disaster strikes.

To be prepared for a system failure, you should have the following:

- Well-documented plans and procedures for recovering from failures when they occur.
- Floppy disks that enable you to restart a computer when you have trouble using the system active or startup volume.
- Documented software and hardware configuration information for your computers.

To reduce system recovery time, it is recommended that you perform the following tasks:

- Put the Windows 2000 Server system and startup and data volumes on separate drives.
- Save the disk configuration data each time you change the configuration by using Disk Management.
- Keep a written record of disk volumes and their sizes.

The remainder of this section describes the disaster protection features of Windows 2000, which you can use to prepare your enterprise for a potential network disaster.

## Creating Backup and Off-Site Storage Policies

Your disaster recovery plan needs to include policies and procedures for backup and restoration of individual computers and entire systems. Your goal is to have clear instructions for recovering your data.

### Backup Policies

When creating backup policies, consider the following strategies:

**Back up all computers or only selected computers.** Do you plan to back up your entire network or will you back up only the servers that have important user files?

**Create network-based backup or local backup.** Will you have a few backup servers with tape devices that read data across the network from all selected servers, or will all users be responsible for backing up their data?

**Use a centralized backup policy or a distributed backup policy.**

Will a single IT group back up all the organization's servers or will each group do its own backup? If the latter is the case, will you establish guidelines for when and how to back up?

Your backup plan needs to include implementation of the following activities:

- Secure both the storage device and the backup media.
- Make copies of all required device drivers so that storage devices can be used in case of disaster. The device drivers are needed to restore operations.
- Create and keep printed copies of backup logs. These are essential for restoring data.
- Keep three copies of the media. Store at least one copy off-site in a properly controlled environment.
- Periodically perform a trial restoration to verify that your backup set can be read and that it contains all the files you want to have backed up. For more information about using specific backup methods and procedures, see “Backup” in the *Microsoft® Windows® 2000 Server Resource Kit Server Operations Guide*.

## Considerations for Off-Site Storage

Consider storage of the following data and information when planning for off-site storage:

- A full backup of the entire system, done on a weekly basis.
- The originals of all installed programs and required device drivers.
- Documents required to facilitate an insurance claim, such as hardware and software inventory records; and copies of purchase orders or receipts for computer hardware and software.
- A copy of information required to reinstall and reconfigure network hardware.

## Creating a Disaster Recovery Plan

To determine what provisions to make for partial or complete loss of data, you need to determine the total cost of rebuilding or replacing the data your organization uses.

Consider the following:

- What are the costs of reconstructing your organization’s financial, personnel, and other business data?
- What does your business insurance cover with regard to replacing lost data?
- How long would it take to reconstruct your business data? How would this translate into lost future business?
- What is the cost per hour of server downtime?

There are several areas that need to be addressed in developing a comprehensive disaster recovery plan. Your plan for data protection needs to answer the following questions:

- What data do you need to back up and how often should you do backups?
- How will you protect critical computer or other hardware configuration information that is not saved during normal backups?
- What data needs to be stored on-site, and how will you physically store it?
- What data needs to be stored off-site, and how will you physically store it?
- What training is required so that server operators and administrators can respond quickly and effectively if an emergency occurs?

Test your plan for recovering and restoring your organization's critical data and keep copies of your disaster recovery plan both on-site and off-site.

## Testing System Recovery Strategies

Testing is an important part of being prepared for disaster recovery. The skill and experience of the administrators and operators is a major factor in getting a failed computer or network back online with minimal cost and disruption to your business. You need IT personnel who are trained in troubleshooting problems and performing system recovery procedures.

Be sure to test recovery procedures before bringing a new server into production.

Testing needs to include:

- Making sure that your Windows 2000 startup disks function correctly.
- Testing your uninterruptible power supply (UPS) on the computers running Windows 2000 Server and on hubs, routers, and other network components.
- Testing your disaster recovery plan.
- Performing full or partial restorations from your daily, weekly, and monthly backup media.

## Practicing Recovery Procedures

You can use testing to try to predict failure situations and to practice recovery procedures. Be sure to do stress testing and test all functionality.

Some of the failures that you need to test include:

- Individual computer components such as hard disks and controllers, processors, and RAM.
- External components such as routers, bridges, switches, cabling, and connectors.

The stress tests that you set up need to include:

- Heavy network loads.
- Heavy disk I/O to the same disk.
- Heavy use of file, print, and applications servers.
- Heavy use by users who log on simultaneously.

## Documenting Recovery Procedures

You need to develop step-by-step procedures for getting a computer or network back online after a disaster. Create an operations handbook that includes the following procedures:

- Performing backups
- Implementing off-site storage policies
- Restoring servers and the network

You should review your documentation when you make configuration changes to your computers or network. Updating the documentation is particularly important when you install new versions of the operating system or change the utilities or tools that you use to maintain your system.

## Planning Task List for Storage Management

Table 19.4 presents a summary checklist that you can use to determine your storage needs and how you will address them.

**Table 19.4 Summary of Storage Planning Tasks**

<b>Task</b>	<b>Location in Chapter</b>
Assess your storage needs.	Improving Your Storage Management Functions
Select a data storage system.	Improving Your Storage Management Functions
Plan for storage management, including removable and remote storage.	Managing Disk Resources
Develop strategies for optimizing storage management.	Optimizing Data Management
Develop strategies for data protection.	Enhancing Data Protection
Develop strategies for backup and disaster recovery.	Improving Your Disaster Recovery Capabilities